arXiv:1607.02437v1 [cs.DS] 8 Jul 2016

Robust Assignments via Ear Decompositions and Randomized Rounding*

David Adjiashvili [a] Viktor Bindewald [b]
Dennis Michaels [b]

[a] Institute for Operations Research, Department of Mathematics, ETH Zurich, Raemistrasse 101, 8092 Zurich (Switzerland)

[b] Department of Mathematics, TU Dortmund University, Vogelpothsweg 87, 44227 Dortmund (Germany)

Many real-life planning problems require making a priori decisions before all parameters of the problem have been revealed. An important special case of such problem arises in scheduling problems, where a set of tasks needs to be assigned to the available set of machines or personnel (resources), in a way that all tasks have assigned resources, and no two tasks share the same resource. In its nominal form, the resulting computational problem becomes the *assignment problem* on general bipartite graphs.

This paper deals with a robust variant of the assignment problem modeling situations where certain edges in the corresponding graph are *vulnerable* and may become unavailable after a solution has been chosen. The goal is to choose a minimum-cost collection of edges such that if any vulnerable edge becomes unavailable, the remaining part of the solution contains an assignment of all tasks.

We present approximation results and hardness proofs for this type of problems, and establish several connections to well-known concepts from matching theory, robust optimization and LP-based techniques.

Keywords. robust optimization ♦ matching theory ♦ ear decomposition ♦ randomized rounding ♦ approximation algorithm

*An extended abstract of this manuscript appeared in I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, D. Sangiorgi (Eds.), Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), July 12-15, 2016, Rome (Italy).

1 Introduction

The need for incorporating system reliability into decision making has sprung wide-spread interest in optimization models which incorporate data uncertainty in the last decades. The latter trend has lead to the development of several new theories including the popular field of *Robust Optimization*. In robust optimization the nominal optimization problem is equipped with a set of *scenarios*, representing various possible states of nature that may occur after the solution to the problem is chosen. The goal is to determine a solution that performs well, in terms of feasibility or cost, in a *worst case* realization of the state of nature.

The *Assignment Problem* is one of the most fundamental optimization problems arising in many reliability-sensitive systems. In its nominal form, the input consists of a collection of n_T *tasks*, a set of n_R *resources* (with $n_T \leq n_R$), and *assignment costs* $c_{i,j}$ representing the cost associated with assigning resource i to task j. The set of allowed assignments can be represented by a bipartite graph $G := (R \dot\cup T, E)$ where each resource i corresponds to a node $r_i \in R$, each task j corresponds to a node $t_j \in T$, and the edge $\{r_i, t_j\}$ is present in E if the j-th task can be assigned to resource i. The goal is to find a minimum-cost *matching* $M \subseteq E$ that covers all nodes in T, i.e. a set of non-adjacent edges that is incident to every node in T. In the following, a subset M satisfying that property is called *an assignment*. A thorough introduction to the assignment problem can be found in the book of Burkard, Dell'Amico, and Martello [11].

This paper deals with a natural robust counterpart of the assignment problem that is defined as follows. An instance of RAP consists of a bipartite graph $G = (R \dot\cup T, E)$ and a non-negative cost vector $c \in \mathbf{R}^E_{\geq 0}$ representing a nominal assignment problem. Furthermore, a collection $\mathcal{F} \subseteq 2^E$ of subsets of edges are given where each $F \in \mathcal{F}$ induces a failure scenario that leads to a deletion of F from G, i.e. if scenario $F \in \mathcal{F}$ emerges, all edges from F are deleted from G. The goal is to find a *redundant assignment* $X \subseteq E$ of minimum cost with the property that, for every $F \in \mathcal{F}$, the set $X \setminus F$ contains an assignment in G (i.e. the graph $(R \dot\cup T, X \setminus F)$ contains an assignment).

The robustness paradigm considered in this paper fits into the concept of *redundancy-based robustness* – a well-motivated and widely studied approach (See Bertsimas, Brown and Caramanis [9], and Herroelen and Leus [36] for an overview of different robustness concepts). Some of the problems falling into

this category include the minimum k-edge connected spanning subgraph problem (Cheriyan, Sebő and Szigeti [18], Gabow et al. [30]) and the robust facility location problem (Jain and Vazirany [39], Swamy and Shmoys [51], Chechik and Peleg [13]).

Adjiashvili, Stiller and Zenklusen [3] introduce a robustness model called *bulk-robustness*, which combines the standard redundancy based robustness approach with a non-uniform failure model. In its general form, a bulk-robust counterpart of a combinatorial optimization problem consists of an instance of the nominal problem, as well as a collection of scenarios, each comprising an arbitrary set of resources that may fail simultaneously. The goal is to determine a minimum-cost set of resources that contains a feasible solution, even when the resources in any single failure scenario become unavailable. In the language of bulk-robustness, RAP is the bulk-robust assignment problem.

It is important to remark that several other robust counterparts of the assignment problem have been considered in the literature under the same, or similar, names. A brief review on relevant models and works existing in the literature is given in Section 2.

In the remainder of this section, a few motivating applications for RAP are provided, some connections to related notions in matching theory are established, and main results as well as technical contributions of this paper are briefly discussed.

1.1 Motivation

The most natural applications of RAP, and redundancy-based robust optimization in general, emerge in situations where resources can not be easily made available on demand. In such applications, any resource intended for deployment at a certain point in time must be *reserved* at an earlier stage, and thus made available for potential deployment. Examples of such applications range from the construction of robust power transmission networks (Hajiaghayi, Immorlica, and Mirrokni [34]) to supply chain management (Tang [52]).

While we think that RAP can be a useful model to incorporate robustness in any assignment model with up-front decisions of the latter type, a few, more concrete, applications are brought hereafter.

Flexible Designs for Manufacturing Processes. Flexible designs for manufacturing processes have attracted significant attention in the operations research community in recent years. The motivation in this topic is the need of manufacturing systems in modern economies to be able to adapt to quickly changing demand patterns. The flexible design is modeled as a problem of selecting a set of edges in a bipartite graph with sides corresponding to plants and products, respectively. An edge between a plant and a product means that the plant can produce the corresponding product. While plants have fixed deterministic capacities, the demand for products can vary, and is assumed to either be random, or materialize from a known uncertainty set. The goal is to choose a cheap set of edges that will maximize the expected, or the worst-case flow between the two sides of the graph. In this view, RAP can be seen as a problem in process flexibility, in which uncertainty lies in the structure of the graph, instead of the demand patters, and the goal is to always satisfy the full demand. For an overview of related results we refer to the papers of Chou et al. [19], Simchi-Levi and Wei [49, 50], and Désir et al. [24].

Staff Training. Large companies often employ intensive training programs for their employees, designed to adapt the available pool of skills to their dynamic needs. For instance, developing new software products often requires improved knowledge on recent technologies that employees have to be trained for. It is natural to incorporate the incurred training costs into the task allocation problem which, in turn, naturally corresponds to an assignment problem. The cost of assigning an employee to perform a given task in the project corresponds to the training cost incurred by training the employee to perform this task.
In a more realistic scenario, some employee to task assignments might become unavailable *even if the employee were trained to perform the task.* That type of vulnerability is very common, and can be caused, e.g. by employee dissatisfaction from his task assignment and by unexpected inability (due to injury or unavailability of equipment, etc.). RAP is a suitable model for deciding on robust training programs for the project, where skill sets of the employees allow for reassignments even if some employee to task assignment becomes unavailable.

Continuity of Service. In industries such as health care and consulting, it is often desirable to maintain very stable client to operator relationships. A typical example is a nursing home where elderly people feel more secure and comfortable if the nurses taking care of them do not change often. This is called

continuity of care in the health care literature (Carello and Lanzarone [12]). In the context of consulting the benefit is the reduction of the loss of undocumented knowledge and employees' stress levels due to decreased number of reassignments. Indirectly this also improves the customer's experience.
Again, the underlying resource allocation problem can be modeled with the help of a bipartite graph $G = (R \dot\cup T, E)$ where nodes from R represent the staff members, while a node from T corresponds to a certain *task*. In this setting, a task can now be any combination of a patient, a shift, and a certain type of service. The edges in G encode whether a staff member is able to perform a task, and the uncertainty set $\mathcal{F} \subseteq 2^E$ contains all edges for which the corresponding staff member has to change a shift with high probability for some reason. From the continuity of service perspective, the aim is to find a smallest subset X of edges in G containing an assignment not using F, for every $F \in \mathcal{F}$. This problem gives rise to an unweighted RAP instance on G and scenario set $\mathcal{F}$.

1.2 Overview of results and techniques

This paper addresses the computational complexity of RAP. The main contributions are approximation algorithms and hardness of approximation results. The study of approximation algorithms is especially justified by showing that RAP is NP-hard, even in very restricted variants.

Notation. To keep notation short, the following standard notation from graph theory is frequently used throughout the paper. Let a graph $G := (R \dot\cup T, E)$ and subsets $H \subseteq E$ and $U \subseteq R \dot\cup T$ be given. The set of all nodes in the graph G covered by edge set H is abbreviated by $V[H]$, and the set of all edges in G having both end nodes in U is denoted by $E[U]$. Furthermore, for a graph G its node set is denoted by $V[G]$ and its edge set by $E[G]$. For two graphs G and $\bar{G}$ we call $\bar{G}$ a subgraph of G if $V[\bar{G}] \subseteq V[G]$ and $E[\bar{G}] \subseteq E[G]$. $G[H] := (V[H], H)$ and $G[U] := (U, E[U])$ refer to the subgraphs of G induced by H and by U, respectively. $G - H$ represents the graph obtained from G by deleting all edges in H, while $G - U$ is used for the graph obtained from G by deleting all nodes of U and all edges incident to some node in U. For two graphs $G' := (R' \dot\cup T', E')$ and $G'' := (R'' \dot\cup T'', E'')$, their union defined as $\big((R' \cup R'') \dot\cup (T' \cup T''),\ E' \cup E''\big)$ is denoted by $G' + G''$.
A subgraph $\bar{G}$ of G is called *spanning* if each node of G is incident to at least one edge of $\bar{G}$.

1.2.1 Problem setting

The assignment problem has a well-known interpretation as a *bipartite matching problem* in the graph $G = (R \dot\cup T, E)$. It is, hence, natural to view RAP as a robust version of the bipartite matching problem, i.e. to find a minimum-cost subset $M \subseteq E$ such that, for every scenario $F \in \mathcal{F}$, the set $M \setminus F$ contains a matching covering all nodes in T. Moreover, if $|R| = |T|$ additionally holds, the problem becomes a robust variant of the *perfect matching* problem.

In the following, this point of view is adopted as it facilitates a clearer exposition of results and highlights an inherent connection between RAP and *matching-covered graphs*, a notion that is repeatedly used in this paper to develop approximation algorithms.

The next statement shows that it suffices to consider RAP on balanced bipartite graphs. It also implies that feasibility conditions on RAP can be stated in terms of perfect matchings.

Proposition 1. *Any* RAP *instance can be efficiently transformed to an equivalent weighted* RAP *instance with a balanced bipartite graph such that, for all $\alpha \geq 1$, any α-approximation for the new instance can be used to efficiently construct an α-approximation of the original instance.*

Proof. Consider any RAP instance on an unbalanced graph $G := (R \dot\cup T, E)$ with $|T| < |R|$, uncertainty set $\mathcal{F}$ and cost vector $c \in \mathbf{R}^E_{\geq 0}$. To transform G into a balanced graph $G' := (R' \dot\cup T', E')$, with $|R'| = |T'|$, a set D of dummy task nodes of cardinality $|D| = |R| - |T|$ is introduced, and each $d \in D$ is connected with every resource node $r \in R$. Further, let E_D be the set of those newly introduced edges, and set $R' := R$, $T' := T \dot\cup D$, $E' := E \dot\cup E_D$ and $\mathcal{F}' := \mathcal{F}$. Finally, choose the new cost vector $c' \in \mathbf{R}^{E'}_{\geq 0}$ such that $c'_e = 0$, if $e \in E_D$, and $c'_e = c_e$, for all $e \in E' \setminus E_D = E$. This procedure can be performed in polynomial time and leads to the desired instance on a balanced graph G'. Now, let $V[E_D]$ denote all nodes from G' covered by E_D. As $T \cap V[E_D] = \emptyset$, there exists a matching $M \subseteq E$ in G that covers T if and only if there is a perfect matching M' in G'. To see this, note that any matching $M \subseteq E$ covering T can be extended to a perfect matching in G' by adding some new edges from E_D while any perfect matching M' in G' must contain a matching $M \subseteq E$ covering T. As the new edges from E_D have zero costs, it follows that $c(M) = c'(M')$. This implies that the transformation preserves quality in terms of approximation. ■

It is important to note that the transformation considered in Proposition 1 produces a *weighted* instance with a balanced bipartite graph, even if the original instance had unit weights. We will hence not be able to use this result in a black-box fashion for our approximation algorithm for the unweighted case. We bring the details in Section 4.1.

The discussion above justifies to focus on RAP defined on balanced bipartite graphs. From now on, it is further assumed that each failure scenario $F \in \mathcal{F}$ is composed of a single edge, i.e. $|F| = 1$. We can thus henceforth assume that $\mathcal{F}$ is simply the set of vulnerable edges, where the scenarios correspond to the failure of any single edge of $\mathcal{F}$. As we show in this paper, this special case of RAP is already interesting both from the application, and the algorithmic points of view.

Summing up, the specific variant of RAP considered in this paper is the following.

Problem (The Robust Assignment Problem (RAP))**.**

- Input: Tuple $(G, \mathcal{F}, c)$, where $G := (R \dot\cup T, E)$ is a balanced, bipartite graph, i.e. $|R| = |T|$, $\mathcal{F} \subseteq E$ is a set of vulnerable edges, and $c \in \mathbf{R}^E_{\geq 0}$ is a non-negative cost vector.
- Output: If exists, an optimal solution for

$$\begin{array}{ll} \min & c(X) \\ \text{s.t.} & \forall f \in \mathcal{F} : X \setminus \{f\} \text{ contains a perfect matching in } G \\ & X \subseteq E. \end{array} \qquad \text{(RAP)}$$

For a given instance of RAP n and m denote the number of nodes and edges of the underlying graph G, i.e. $n := |R| + |T|$ and $m := |E|$.

Two special cases of RAP are of particular interest. The first case, denoted by card-RAP, is given when the cost function of RAP is *unweighted*, i.e. when the task is to find a feasible (robust) solution of minimum cardinality. In the second case, every edge in the underlying graph $G = (R \dot\cup T, E)$ is assumed to be vulnerable, i.e. $\mathcal{F} = E$. In this case, the RAP instance is called *uniform*.

Before we move on, we briefly treat the problem of deciding whether a RAP instance is feasible. Observe that, for an arbitrary instance $\mathcal{I} = (G, \mathcal{F}, c)$ of RAP, the feasible set is *monotonic* in the sense that any superset of a feasible solution is feasible as well. Thus, $\mathcal{I}$ is a feasible instance if and only if the edge set E of

G is a feasible solution, or equivalently, $E \setminus \{f\}$ contains a perfect matching in G, for every $f \in \mathcal{F}$. The latter condition can be checked using any polynomial algorithm for finding maximum matchings in bipartite graphs. Therefore, feasibility of RAP can be efficiently verified. We will henceforth assume that any RAP instance considered in this paper is feasible.

It is worth to remark that the latter is no longer true when the uncertainty set $\mathcal{F}$ is given implicitly. For example, consider a balanced bipartite graph $G := (R \dot{\cup} T, E)$ and an uncertainty set $\mathcal{F} := \{F \subseteq E \colon |F| = k\}$, for some $k \in \mathbf{Z}_{>1}$, presented by the parameter k. Then, checking for every $F \in \mathcal{F}$ whether $E \setminus F$ contains a perfect matching in G is equivalent to the problem of deciding if the so-called *matching preclusion number* of G is at most k. The latter problem was proved to be NP-complete for bipartite graphs by Dourado et al. in [26, Thm. 2].

1.2.2 Matching-Covered Graphs

The algorithmic results derived in this paper rely on a tight connection between RAP and *matching-covered* graphs, a well-known notion in matching theory. Recap that a graph is *matching-covered* if each of its edges appears in some perfect matching [1].

It turns out that inclusion-wise minimal solutions of any RAP instance are matching-covered as the following proposition states.

Proposition 2. *Let $\mathcal{I} := (G, \mathcal{F}, c)$ be any feasible* RAP *instance. Then, a subset X of edges in G is an inclusion-wise minimal feasible solution to $\mathcal{I}$ if and only if its induced subgraph $G[X]$ spans G and is inclusion-wise minimal with the properties of being matching-covered and that every isolated edge $e \in G[X]$ is not vulnerable, i.e. $e \notin \mathcal{F}$.*

Proof. "only if" part. Let X be an inclusion-wise minimal feasible solution to $\mathcal{I}$, and let $e \in X$ be any edge. Firstly, if e does not appear in any perfect matching, $X \setminus \{e\}$ is feasible to $\mathcal{I}$. Thus, inclusion-wise minimality of X implies that $G[X]$ is matching-covered. Secondly, if $f \in \mathcal{F}$ is an isolated edge in $G[X]$ then $X \setminus \{f\}$ cannot contain a perfect matching in G. This contradicts that X is feasible. Thirdly, assume that $G[X]$ contains a proper spanning subgraph

[1] The notion of matching-covered graphs is originally introduced for connected graphs. In this paper we use this term also for disconnected graphs. Note further that some authors use synonymously the notion 1-*extendable* or, in the bipartite case, *elementary* (cf. [42]).

$G[X']$, induced by some $X' \subsetneq X$, that is matching-covered and that has no isolated edges from $\mathcal{F}$. Consider an arbitrary $f \in X' \cap \mathcal{F}$. As f is not isolated in $G[X']$, there is an $e' \in X'$ adjacent to f. Since $G[X']$ spans G and is matching-covered, there is a perfect matching in G that contains e' and that does not contain f. This, however, shows that X' is also feasible to $\mathcal{I}$ contradicting the inclusion-wise minimality of X.

"if" part. Let X be a subset of edges from G such that its induced subgraph $G[X]$ spans G and is inclusion-wise minimal with respect to the properties of being matching-covered and that every isolated edge in $G[X]$ is not vulnerable. Showing that X is feasible to $\mathcal{I}$ is similar to the proof of the feasibility of X' in the first direction of the proof. Assume now that there is an $\bar{X} \subseteq X$, $\bar{X} \neq X$, that is feasible to $\mathcal{I}$. This firstly implies that the induced subgraph $G[\bar{X}]$ spans G. Secondly, no vulnerable edge $f \in \bar{X} \cap \mathcal{F}$, if exists, can form an isolated edge in $G[\bar{X}]$. Moreover, it can be assumed that $G[\bar{X}]$ is also matching-covered, as otherwise, each $e \in \bar{X}$ not extendable to a perfect matching in $G[\bar{X}]$ can be removed from $\bar{X}$. This way, a proper subset of X is obtained whose induced subgraph spans G and is matching-covered. This concludes the proof. ■

Proposition 2 provides a very useful characterization of inclusion-wise minimal solutions of RAP, as it allows to make use of various results on matching-covered graphs to design algorithms for RAP. In particular, it allows to identify feasible subgraphs and augment them to feasible solutions for the entire instance by adding structures that maintain the property of being matching-covered.

1.2.3 Results for RAP

One important contribution of this paper is to show that approximating RAP is as hard as to approximate the well-known Set Cover Problem. This is stated in the following theorem.

Theorem 3. RAP *admits no polynomial* $d \log n$*-approximation algorithm for any* $d < 1$*, unless* NP $\not\subseteq$ DTIME$(n^{\log \log n})$[2]*. This is true even for uniform* RAP*.*

[2]Recap that NP $\subseteq$ DTIME$(n^{\log \log n})$ would imply the existence of quasi-polynomial time algorithms for NP-hard problem.

Theorem 3 motivates developing an approximation algorithm for RAP with the matching asymptotic bound $O(\log n)$, imposed by the previous theorem. We achieve this goal in the next theorem.

Theorem 4. RAP *admits a randomized polynomial* $O(\log n)$*-approximation algorithm.*

For the proof of Theorem 4, an algorithm with the desired approximation quality is presented. The algorithm constructs iteratively a solution maintaining the invariant that, at any iteration, the edges selected so far, form a matching-covered graph. It is, however, unclear how to arrive at the desired approximation for RAP when only properties of matching-covered graphs are taken into account. Therefore, the latter technique is combined with additional tools from the theory of linear programming (LP) and randomized rounding. More precisely, the algorithm starts with solving an LP relaxation of RAP, derived from a natural integer linear programming (ILP) formulation of the problem. The fractional solution obtained this way is used to guide an iterative randomized procedure. In each iteration a fractional bipartite matching corresponding to part of the fractional solution is selected. A decomposition of this fractional matching into a convex combination of integral matchings is then used to randomly pick one matching, and a carefully selected subset of this matching is added to the current solution. To bound the quality, it does not suffice to bound the number of iterations, or the expected number of times an edge is part of a candidate matching. Instead, a discharging argument, that assigns costs to nodes depending on the graph selected so far, is used.

A detailed presentation of all results for RAP, including the proofs, is given in Section 3.

1.2.4 Results for card-RAP

Our main complexity result for the unweighted RAP states that it is NP-hard to approximate within some constant $\delta > 1$, even for uniform uncertainty sets.

Theorem 5. *For some constant* $\delta > 1$*, there is no polynomial* δ*-approximation for uniform* card-RAP*, unless* P= NP*.*

On the positive side, we are able to use the strong connection between RAP and matching-covered graphs to develop a constant factor approximation algorithm for card-RAP using so-called *ear decompositions*.

Theorem 6. card-RAP *admits a polynomial* 1.5*-approximation algorithm in the uniform case, and a* 3*-approximation algorithm in the non-uniform case.*

The algorithm starts by producing an ear decomposition of the input graph. Then, it selects a certain subset of the edges to be part of the solution, by processing the ears in the decomposition in the order given by the decomposition, and omitting the edges corresponding to ears of length one.

Theorems 5 and 6 imply that, assuming P $\neq$ NP, the true approximability thresholds for uniform card-RAP and card-RAP lie in the intervals $[\delta, 1.5]$ and $[\delta', 3]$, for some $\delta, \delta' > 1$. These results are proved in Section 4.

To complete the complexity landscape of card-RAP, the case with only two vulnerable edges is also considered. This special case comprises the simplest possible variant of RAP that is not equivalent to a nominal assignment problem [3]. It turns out that this special case is already NP-hard, thus drawing a sharp threshold for polynomial solvability of RAP.

Theorem 7. card-RAP *is* NP*-hard even when restricted to instances with two vulnerable edges, i.e. with* $\mathcal{F} = \{f_1, f_2\}$.

The proof of Theorem 7 relies on a connection to a new NP-hard problem called Shortest Nice Path Problem (SNPP). The goal is to partition a graph into a path connecting two given nodes and a matching such that their union covers all nodes, so as to minimize the length of the path. This problem might be interesting in its own right. The formal definition of SNPP and the proof of Theorem 7 is presented in Section 5.

Remark 8. To the best of the authors' knowledge, this is the first example of an NP-hard robust counterpart of a polynomial optimization problem, with a constant number of vulnerable resources. Note that there are many examples of optimization problems that become NP-hard when the robust counterpart is allowed to contain a constant number of scenarios (see e.g. [40]). However, in all such examples, each scenario affects a *non-constant* number of resources.

[3] Observe that the case of a single vulnerable edge $\mathcal{F} = \{f\}$ is solvable by reporting any minimum-cost perfect matching in the graph $(R \dot\cup T, E \setminus \{f\})$ as a solution.

2 Related work

Redundancy-based robustness is a paradigm that motivates many well-studied problems, including the minimum k-connected subgraph problem (see Gabow et al. [30], Cheriyan, Sebő and Szigeti [18], and Sebő and Vygen [48]), survivable and robust network design problems (Jain [38], Chekuri [16], Adjiashvili, Stiller and Zenklusen [3], and Adjiashvili [2]), robust facility location problems (Jain and Vazirani [39], and Smamy and Shmoys [51]), robust spanner problems (Chechik et al. [14], and Dinitz and Krauthgamer [25]), and many more. All of the latter models bare a close resemblance to RAP: they assume resources to be vulnerable and ask to find a minimum-cost set of resources that contains a desired structure even in case any vulnerable resource, or set of resources, fails.

A relatively new approach to redundancy-based robustness is the incorporation of non-uniform uncertainty sets ([3, 2]). RAP is seen as a robust model of this type, as both vulnerable and invulnerable edges are allowed to appear in the same instance.

The study of robustness with respect to cost uncertainty is initiated by Kouvelis and Yu [40], and Yu and Yang [56]. These works mainly consider the min-max model, where the goal is to find a solution that minimizes the worst-case cost according to the given set of cost functions. A survey on this topic is given by Aissi, Bazgan and Vanderpooten [5]. A closely related class of multi-budgeted problems is received considerable attention (see e.g. Grandoni et al. [32] and references therein). The latter work includes variants of the related multi-objective matching problem for which a polynomial time approximation scheme is presented.

Various variants of the robust matching problems are considered in the literature. Hassin and Rubinstein [35], and Fujita, Kobayashi, and Makino [28] study the following notion of an α-robust matching. A perfect matching M in a weighted graph is α-robust (for $\alpha \in (0, 1]$), if for every $p \leq |M|$, the p heaviest edges of the matching have total weight at least α times the weight of a maximum-weight matching of size p. In [35] the authors prove that the complete graph K_n contains a $\frac{1}{\sqrt{2}}$-robust matching and this bound is tight in general. Additionally, the authors provide a polynomial-time algorithm to find such a matching. Building upon these results Fujita, Kobayashi, and Makino [28] prove that the problem of deciding if the input graph has α-robust matching with $\alpha \in (\frac{1}{\sqrt{2}}, 1)$ is NP-complete, and extend the original algorithm to the matroid intersection problem.

Deineko and Woeginger [23] show that the min-max-robust assignment problem with a fixed number of scenarios is equivalent to the *exact perfect matching problem*, a famous problem with unknown complexity status. In the case of a variable number of scenarios the min-max-robust problem is NP-hard, as proved by Aissi, Bazgan and Vanderpooten [4].

Laroche et al. [41] investigate the following robust variant of the matching problem. Find the maximum number of nodes that can be removed from the larger side of the bipartition such that all nodes from the opposite bipartition can still be matched. The authors also provide a polynomial-time algorithm for this problem.

Bertsimas, Brown and Caramanis [9] provide a comprehensive survey of the different facets of robust optimization. Surveys on robust combinatorial optimization are given in the PhD theses of Olver [43] and Adjiashvili [1].

Brigham et al. [10] study the minimum number of edges to be removed from a graph to arrive at a graph without a perfect matching, i.e. the matching preclusion number. For several graph classes the matching preclusion number as well as all optimal solutions are computed. Cheng et al. [17] extend this concept by excluding the obvious case when the edges' removal produces isolated nodes introducing the *conditional matching preclusion sets*. For several basic graph classes optimal such sets are presented. Dourado et al. [26] study the closely related problem of finding a robust and recoverable matching: Given a graph G, a set F of edges, and two integers r and s, does the graph G have a perfect matching M such that, for every choice F' of r edges from F, the graph $G - F'$ contains a perfect matching M' having a symmetric difference with M of size at most s? Such a matching M is called *r-robust* and *s-recoverable*. The authors prove hardness of several related problems and present some tractable cases.

Chegireddy and Hamacher [15] investigate the problem of finding the k best perfect matchings in terms of weight in a given graph and provide an $O(kn^3)$ algorithm.

Hung, Hsu and Sung [37] consider the so-called *most vital edges* with respect to a weighted bipartite matching, i.e. the edges causing the largest decrement in the value of the maximum matching upon removal and provided an $O(n^3)$ algorithm. Zenklusen [57] analyzes the related interdiction problem, i.e. the problem of choosing a set R of edges (or nodes) in some weighted graph G respecting a given budget in order to minimize the weight of maximum matching in subgraph of G resulting by removal of R. Several special cases of that

problem are shown to be NP-hard. An $O(1)$-approximation algorithm is also provided.

Darmann et al. [21] study the hardness of the maximum matching problem with the additional structure of conflict (respectively, forcing) graphs. Such graphs describe pairs of edges such that at most (respectively, at least) one of the two is to be part of the solution. The problem is shown to be NP-hard even for very restricted classes of conflict and forcing graphs. Öncan, Zhang and Punnen [44] presents further complexity results and heuristics for the perfect matching problem with conflict constraints.

Structural and graph-theoretic aspects of modifications of the matching problem have also been considered in the literature.

Plesník [45] proves that an $(r-1)$-edge-connected r-regular graph remains perfectly matchable after removing $r-1$ arbitrary edges. In the context of RAP this means that a 2-regular edge-connected subgraph of a (non-)bipartite graph is feasible in the uniform case.

Some generalizations of matching-covered graphs, such as *n-extendable* graphs, were investigated in the literature. A graph is called *n-extendable* if any matching of size n can be extended to a perfect matching. Thus, matching-covered graphs form the class of 1-extendable graphs. Moreover, n-extendability can be seen as a special case of the so-called *$E(m,n)$–property*. A graph satisfies the *$E(m,n)$–property* if, for any pair (M,N) of disjoint matchings with $|M| = m$ and $|N| = n$, the graph has a perfect matching that contains N and does not contain any element of M. Such graphs were studied by several authors, e.g. Porteous and Aldred [46], Aldred et al. [7], and Aldred and Plummer [6].

Wang, Yuan and Zhou [55] considered k-edge-deletable IM-extendable graphs that are characterized as follows. A graph G has the latter property if after removing any set F of k edges from it, every induced matching M of $G-F$ (i.e. a matching such that $E[V[M]] = M$) is contained in a perfect matching of $G-F$.

Several publications deal with algorithms to identify all edges of a graph being part of some maximum matching. Regin [47], Tassa [53], and Costa [20] provided results on bipartite graphs, while Carvalho and Cheriyan [22] considered general graphs. For bipartite graphs, Valencia and Vargas [54] studied the weighted version of this problem that asks to find all edges belonging to some minimum weight perfect matching.

3 Hardness and approximability of RAP

This section deals with the approximability of RAP. In Subsection 3.1, it is first shown that the well-known *Set Cover Problem* can be polynomially reduced to weighted RAP. From that reduction, it is then concluded that under a mild and widely accepted assumption the bound on a reachable approximation guarantee on RAP is of order $\log n$, proving Theorem 3. Secondly, an asymptotically tight randomized algorithm is developed in Subsection 3.2. The existence of this algorithm implies the correctness of Theorem 4.

3.1 Reduction from the Set Cover Problem

The proof of the hardness of the approximability of RAP relies on a reduction from the Set Cover Problem[4].

Problem (Set Cover Problem (SCP))**.**

- Input: Tuple $([k], \mathcal{S})$, where $[k] = \{1, \ldots, k\}$ is a finite ground set and $\mathcal{S} := \{S_1, \ldots, S_l\}$ is a collection of subsets of $[k]$, for some $k, l \in \mathbf{Z}_{\geq 1}$.
- Output: A collection $\mathcal{C} \subseteq \{S_1, \ldots, S_l\}$ of minimum cardinality such that that $\bigcup_{S \in \mathcal{C}} S = [k]$ holds.

For a given instance $([k], \mathcal{S})$, the existence of any cover for ground $[k]$ can efficiently be verified, simply by checking if $\bigcup_{S \in \mathcal{S}} S = [k]$ holds. From now on, it is hence assumed that any considered SCP instance contains at least one feasible cover solution.

For any feasible SCP instance $([k], \mathcal{S})$ consider now the balanced bipartite graph $G := (R \dot{\cup} T, E)$ illustrated in Figure 1 and constructed as follows.

(T1) For each $s \in [k]$, a node u_s is introduced and added to T, and for each $S \in \mathcal{S}$, a node v_S is introduced and added to R. Furthermore, G contains the edge $\{u_s, v_S\}$ if and only if $s \in S$. All these edges form subclass E_2.

(T2) For each $s \in [k]$, a copy $\bar{u}_s$ of node u_s is introduced and added to R, and u_s and $\bar{u}_s$ are connected by an edge. These edges are composited to E_1.

[4]For the underlying NP-complete decision problem, see [31, Problem SP5].

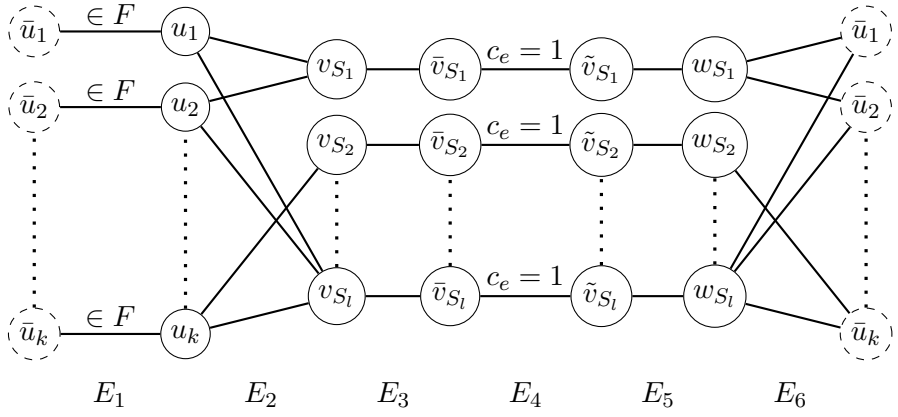

Figure 1: The figure shows the graph G that is constructed from an instance $([k], \{S_1, \ldots, S_l\})$ of SCP. For the purposes of legibility, the dashed nodes representing node set $\{\bar{u}_s \mid s \in [k]\}$ appear twice in the figure.

(T3) For each $S \in \mathcal{S}$, two copies $\bar{v}_S$ and $\tilde{v}_S$ of node v_S are introduced whereupon the first copy is added to T, and the second one is added to R. The original node v_S is connected with $\bar{v}_S$, and $\bar{v}_S$ is further connected with $\tilde{v}_S$. All edges of the form $\{v_S, \bar{v}_S\}$ yield subclass E_3, while all edges of type $\{\bar{v}_S, \tilde{v}_S\}$ belong to subclass E_4.

(T4) For each $S \in \mathcal{S}$, a fourth type of node is introduced that is denoted by w_S and added to T. Each node w_S is connected with the corresponding node $\tilde{v}_S$. These edges form subclass E_5. Finally, an edge between w_S and $\bar{u}_s$ is introduced and added to subclass E_6 if and only if $s \in S$.

With some $\mathcal{F} \subseteq E$ and some $c \in \mathbf{R}^E_{\geq 0}$, the tuple $(G, \mathcal{F}, c)$ then defines an instance of RAP. Note that both E_2 and E_6 encode whether element $s \in [k]$ is contained in subset $S \in \mathcal{S}$, or not. E_3 and E_5 ensure the feasibility of the RAP instance, while the edges in E_4 are used to indicate which elements from $\mathcal{S}$ are chosen to cover the ground set $[k]$.

The next lemma highlights the close relation between SCP and the RAP.

Lemma 9. *Let $\mathcal{I} := ([k], \mathcal{S})$ be an instance of* SCP*, and let $\mathcal{I}' := (G, \mathcal{F}, c)$ be the* RAP *instance with $G := (R \dot{\cup} T, E)$ as constructed by applying steps* **(T1)** – **(T4)***, uncertainty set $\mathcal{F} = E_1$ and cost vector $c \in \mathbf{R}^E_{\geq 0}$ with $c_e = 1$, if $e \in E_4$, and $c_e = 0$, if $e \in E \setminus E_4$.*

(a) $\mathcal{I}'$ can be determined in time polynomially bounded in the input size of $\mathcal{I}$.

(b) *Let $X \subseteq E$ be with $E \setminus E_4 \subseteq X$. Then, X is feasible to $\mathcal{I}'$ if and only if $\mathcal{C}_X := \{S \in \mathcal{S} \mid \{\bar{v}_S, \tilde{v}_S\} \in X\}$ is feasible to $\mathcal{I}$.*

(c) *$\mathcal{I}$ has a cover of size* VAL *if and only if $\mathcal{I}'$ contains a feasible solution $X \subseteq \bar{E}$ with objective value $c(X) =$* VAL.

Proof. (a) Let $\mathcal{I} := ([k], \mathcal{S})$ be given. From the construction of G, it follows that $|T| = |R| = 2|\mathcal{S}| + |k|$ and that $|E| = 3|\mathcal{S}| + |k| + 2\sum_{S\in\mathcal{S}} |S|$ holds, i.e. the size of G is polynomially bounded by the size of $\mathcal{I}$. Further, E_1 and c are, by definition, polynomially bounded in the input size of G. Thus, $\mathcal{I}' = (G, E_1, c)$ can be determined in polynomial time w.r.t. the input size of $\mathcal{I}$.

(b) *"only if" part.* Let X be any feasible solution to $\mathcal{I}'$ with $E \setminus E_4 \subseteq X$, and let $s \in [k]$. To show that s is contained in some set of $\mathcal{C}_X$, consider the edge $f_s = \{\bar{u}_s, u_s\} \in \mathcal{F} = E_1$. As X is feasible to $\mathcal{I}'$, there is a perfect matching $M \subseteq X$ in G with $f_s \notin M$. Since $f_s \notin M$, node u_s must be matched with some node from $\{v_{S_1}, \ldots, v_{S_l}\}$ by the corresponding edge from E_2, i.e. $\{u_s, v_S\} \in X$, for some $S \in \mathcal{S}$ with $s \in S$. Edge $\{u_s, v_S\}$ also covers node v_S. Thus, the node $\bar{v}_S$ must be matched with $\tilde{v}_S$. This implies that $\{\bar{v}_S, \tilde{v}_S\} \in X$ and, hence, that $S \in \mathcal{C}_X$. It follows that s is covered, and that $\mathcal{C}_X$ is a feasible cover for $\mathcal{I}$.

"if" part. Let $\mathcal{C} \subseteq \mathcal{S}$ be a feasible cover for $\mathcal{I}$. Then, define

$$X := (E \setminus E_4) \bigcup \big\{ \{\bar{v}_S, \tilde{v}_S\} \mid S \in \mathcal{C}\big\}, \text{ implying that } \mathcal{C} = \mathcal{C}_X \text{ holds.}$$

Recap further that X is feasible to $\mathcal{I}'$ if and only if $X \setminus \{f_s\}$ contains a perfect matching of G, for all $f_s = \{\bar{u}_s, u_s\} \in \mathcal{F} = E_1$.
To verify the correctness of the latter condition, let $f_s \in \mathcal{F}$ be arbitrary. Note also that $M_0 := E_1 \cup E_3 \cup E_5$ is a perfect matching in G with $f_s \in M_0 \subseteq X$. As $\mathcal{C}$ is a feasible cover for $\mathcal{I}$, there exists a set $S \in \mathcal{S}$ with $s \in S$ and $\{\bar{v}_S, \tilde{v}_S\} \in X$. Now consider the cycle C_0 defined on node set $\{\bar{u}_s, u_s, v_S, \bar{v}_S, \tilde{v}_S, w_S\}$. Observe that all edges in C_0 are contained in X. In particular, it holds that $\{\bar{u}_s, u_s\}, \{v_S, \bar{v}_S\}, \{\tilde{v}_S, w_S\} \in M_0$ and that $\{u_s, v_S\}, \{\bar{v}_S, \tilde{v}_S\}, \{w_S, \bar{u}_s\} \in X \setminus M_0$, i.e. C_0 forms an M_0-alternating cycle in X. Therefore, the set M can be expressed as the symmetric difference $M_0 \triangle C_0$ and is given as

$$\Big(M_0 \setminus \big\{\{\bar{u}_s, u_s\}, \{v_S, \bar{v}_S\}, \{\tilde{v}_S, w_S\}\big\}\Big) \cup \big\{\{u_s, v_S\}, \{\bar{v}_S, \tilde{v}_S\}, \{w_S, \bar{u}_s\}\big\}.$$

Hence, M is a perfect matching in G with $\{\bar{u}_s, u_s\} = f_s \notin M \subseteq X$.

(c) This follows directly from part (b) and by the choice of c.

■

Lemma 9 already implies that any SCP instance can be transformed, in polynomial time, into an equivalent RAP instance. It can further be concluded that RAP is as hard to approximate as SCP. The proof of Theorem 3, however, requires a stronger result, namely that SCP can be even equivalently transformed into an *uniform* RAP instance.

Again, let $\mathcal{I} := ([k], \mathcal{S})$ be any SCP instance, and let $G := (R \dot\cup T, E)$, with $E := \bigcup_{i=1}^{6} E_i$, be the graph constructed from $\mathcal{I}$ by applying the transformation steps **(T1)**–**(T4)** as described above. In general, the uniform version of the RAP instance $\mathcal{I}'$ on G is not equivalent to $\mathcal{I}$. To see this, pick any edge $e = \{\tilde{v}_S, w_S\} \in E_5$ and assume that e is vulnerable, i.e. e belongs to $\mathcal{F}$. Then, it immediately follows that the (unique) edge $e' := \{\bar{v}_S, \tilde{v}_S\}$ in E_4, incident with e, must be part of any feasible solution, as otherwise $\tilde{v}_s$ cannot be matched without using e. Therefore, the only subset $X \subseteq E$ that contains $E \setminus E_4$ and that is feasible to $\mathcal{I}'$ is given by $X = E$. In the light of Lemma 9, this shows that a uniform RAP instance on G has precisely one feasible solution that only represents the trivial feasible cover $\mathcal{S}$. The same line of arguments also works when e is chosen from E_3 instead from E_5.

To overcome the described problem, the graph G must be changed in an appropriate way. For this, the following, additional transformation step is performed on G.

(T5) Each edge $e = \{v, w\} \in E_3 \cup E_5$ is replaced by a cycle of length six. That means, for each e introduce four auxiliary nodes $x_1^{(e)}, x_2^{(e)}, y_1^{(e)}, y_2^{(e)}$. Then, remove each $e = \{v, w\} \in E_3 \cup E_5$ from G, and add the edges $\{v, x_1^{(e)}\}, \{x_1^{(e)}, x_2^{(e)}\}, \{x_2^{(e)}, w\}$ as well as $\{v, y_1^{(e)}\}, \{y_1^{(e)}, y_2^{(e)}\}, \{y_2^{(e)}, w\}$ to G (see Figure 2, for an illustration).

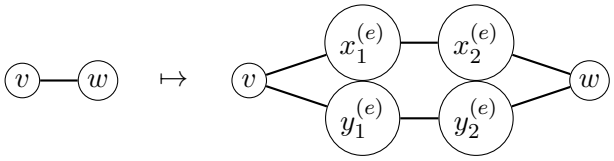

Figure 2: The Figure shows the substitution of an edge $e = \{v, w\}$ by two parallel paths of length three with the help of four auxiliary nodes.

Transformation step **(T5)** yields a new graph $\bar{G} := (\bar{R} \dot{\cup} \bar{T}, \bar{E})$ with $\bar{E} = \bigcup_{j=1}^{6} \bar{E}_j$ and

- $\bar{R} := R \cup \{x_2^{(e)}, y_2^{(e)} \mid e \in E_3 \cup E_5\}, \quad \bar{T} := T \cup \{x_1^{(e)}, y_1^{(e)} \mid e \in E_3 \cup E_5\},$
- $\bar{E}_3 := \big\{ \{v, x_1^{(e)}\}, \{x_1^{(e)}, x_2^{(e)}\}, \{x_2^{(e)}, w\}, \{v, y_1^{(e)}\}, \{y_1^{(e)}, y_2^{(e)}\}, \{y_2^{(e)}, w\} \mid e = \{v, w\} \in E_3 \big\},$
- $\bar{E}_5 := \big\{ \{v, x_1^{(e)}\}, \{x_1^{(e)}, x_2^{(e)}\}, \{x_2^{(e)}, w\}, \{v, y_1^{(e)}\}, \{y_1^{(e)}, y_2^{(e)}\}, \{y_2^{(e)}, w\} \mid e = \{v, w\} \in E_5 \big\},$
- $\bar{E}_1 := E_1$, $\bar{E}_2 := E_2$, $\bar{E}_4 := E_4$, and $\bar{E}_6 := E_6$.

Note that $\bar{G}$ is constructable in polynomial time (in terms of the input size of $\mathcal{I}$). The next lemma shows that SCP can be equivalently restated as a uniform RAP instance on $\bar{G}$.

Lemma 10. *Let $\mathcal{I} := ([k], \mathcal{S})$ be an instance of* SCP*, and let $\bar{\mathcal{I}} := (\bar{G}, \bar{E}, \bar{c})$ be the uniform* RAP *instance with $\bar{G} := (\bar{R} \dot{\cup} \bar{T}, \bar{E})$ constructed by applying steps* **(T1)** – **(T5)***, and with $\bar{c} \in \mathbf{R}_{\geq 0}^{\bar{E}}$, where $\bar{c}_e = 1$, if $e \in \bar{E}_4$, and $\bar{c}_e = 0$, otherwise.*

(a) $\bar{\mathcal{I}}$ can be determined in time polynomially bounded in the input size of $\mathcal{I}$.

(b) Let $X \subseteq \bar{E}$ such that $\bar{E} \setminus \bar{E}_4 \subseteq X$ holds. Then, X is feasible to $\bar{\mathcal{I}}$ if and only if $\mathcal{C}_X := \{S \in \mathcal{S} \mid \{\bar{v}_S, \tilde{v}_S\} \in X\}$ is feasible to $\mathcal{I}$.

(c) $\mathcal{I}$ has a cover of size VAL *if and only if $\bar{\mathcal{I}}$ contains a feasible solution $X \subseteq \bar{E}$ with objective value $\bar{c}(X) =$* VAL*.*

The proof of Lemma 10 is very similar to the proof of Lemma 9, and is hence omitted here. With Lemma 10 at hand, Theorem 3 can be proved.

Proof of Theorem 3. By Lemma 10, any SCP instance $\mathcal{I} := ([k], \mathcal{S})$ can be reduced, in polynomial time, into a uniform RAP instance $\bar{\mathcal{I}} := (\bar{G}, \bar{E}, \bar{c})$, with $\bar{G} = (\bar{R} \dot{\cup} \bar{T}, \bar{E})$ and $\bar{c} \in \mathbf{R}_{\geq 0}^{\bar{E}}$ as specified in the statement of Lemma 10. In particular, it follows from part (c) that a feasible cover $\mathcal{C} \subseteq \mathcal{S}$ in $\mathcal{I}$ has a cardinality of VAL $\in \mathbf{Z}_{\geq 0}$ if and only if there is a feasible solution $X \subseteq \bar{E}$ for $\bar{\mathcal{I}}$ containing VAL different edges out of $\big\{\{\bar{v}_S, \tilde{v}_S\} \mid S \in \mathcal{S}\big\} (= \bar{E}_4 = E_4)$.
Feige [27] showed that, for any $d < 1$, SCP admits no polynomial time $d \log n$-approximation algorithm unless NP $\subseteq$ DTIME$(n^{\log \log n})$. His result also holds when $|\mathcal{S}|$ is polynomial in k, implying that $\bar{G}$ has size polynomial in k.
Combining Feige's result with Lemma 10 completes the proof. ■

3.2 $O(\log n)$-Approximation for RAP

This section presents a polynomial $O(\log n)$-approximation algorithm for RAP, thus proving Theorem 4. After introducing some basic notation, the approximation algorithm is first described for the uniform case. Afterwards, it is explained how the algorithm can be extended to the non-uniform case.

Again, it is assumed that the RAP instance, consisting of a balanced bipartite graph $G = (R \dot\cup T, E)$, an uncertainty set $\mathcal{F} \subseteq E$ and of an non-negative cost vector $c \in \mathbf{R}^E_{\geq 0}$, is feasible. Furthermore, for each $S \subseteq E$, $\chi^S \in \{0,1\}^E$ represents the *incidence vector* of S, i.e. the vector with $\chi^S_e = 1$ if $e \in S$ and $\chi^S_e = 0$, otherwise.

The algorithm is based on an LP-rounding procedure that works with a relaxation of the integer linear formulation of RAP. For this, let $P_G \subseteq \mathbf{R}^E$ be the perfect matching polytope associated with G, i.e. P_G is the convex hull of all incidence vectors of perfect matchings in G.

A standard ILP formulation of RAP involves the following variables.

(i) $x^{-f} \in \{0,1\}^E$ representing a perfect matching in $G - \{f\}$, for all $f \in \mathcal{F}$,

(ii) $y \in \{0,1\}^E$ encoding a feasible solution to RAP.

Then, RAP can be modeled as an ILP as follows.

$$\begin{array}{lrcll} \min & c^\top y & & & \\ \text{s.t.} & x^{-f} & \in & P_G \cap \{x \in \mathbf{R}^E \mid x_f = 0\}, & \text{for each } f \in \mathcal{F}, \\ & y & \geq & x^{-f}, & \text{for each } f \in \mathcal{F}, \\ & x^{-f} & \in & \{0,1\}^E, & \text{for each } f \in \mathcal{F}, \\ & y & \in & \{0,1\}^E. & \end{array} \tag{ILP}$$

The LP-relaxation (LP) is obtained by relaxing all integrality constraints in (ILP). To keep notation short, let $x \in \left(\mathbf{R}^E\right)^E$ be the vector with parts x^{-f}, $f \in \mathcal{F}$. It is straightforward to verify that integer solutions to (ILP) coincide with feasible solutions to the RAP instance.

Now consider the uniform case, i.e. $\mathcal{F} = E$. Let (x, y) be a fractional solution to the corresponding (LP), and select some edge $f \in \mathcal{F}$. Since x^{-f} is contained in $P_G \cap \{x \in \mathbf{R}^E \mid x_f = 0\}$, there exist positive scalars $\lambda_1^{-f}, \cdots, \lambda_k^{-f}$ with $\sum_{i\in[k]} \lambda_i^{-f} = 1$, and perfect matchings $M_1^{-f}, \cdots, M_k^{-f}$ in $G - \{f\}$ such that $x^{-f} = \sum_{i\in[k]} \lambda_i^{-f} \chi^{M_i^{-f}}$. By Caratheodory's theorem, there is a decomposition

of the latter type with k bounded by $O(m) = O(n^2)$. Furthermore given x^{-f}, such a decomposition can be computed in polynomial time using polyhedral techniques [33, Thm. 6.5.11].

Algorithm 1 : Randomized $O(\log n)$-Approximation for RAP

Require: $G = (R \dot\cup T, E)$ with $|R| = |T|$, and $c \in \mathbf{R}^E_{\geq 0}$.
Ensure: A feasible solution X to RAP on G with $\mathcal{F} = E$ and cost vector c.

```
Solve (LP) to obtain an optimal solution (x, y)
X ← ∅
while X is infeasible do
    Select an edge f ∈ F such that X \ {f} contains no perfect matching
    Compute a decomposition of x^{-f} as x^{-f} = Σ_{i=1}^{k} λ_i^{-f} χ^{M_i^{-f}} and select
    one matching M̄ ∈ {M_i^{-f} | i ∈ [k]} with Pr[M̄ = M_i^{-f}] = λ_i^{-f} for all
    i ∈ [k]
    Add to X all edges from M̄ that connect distinct connected components
    in (R ∪̇ T, X)
end while
return X
```

The algorithm performs several iterations of randomized rounding that are based on the latter decomposition of fractional matchings. More precisely, at each iteration, an infeasible set $X \subseteq E$ of edges, that was chosen so far, is augmented with an additional set M of edges chosen randomly as follows. First, an arbitrary edge f is chosen from E among all edges not yet covered by X, i.e. among all $e' \in E$ such that the edge set X selected so far contain no perfect matching that does not include e'. Next, a decomposition of the vector x^{-f} as a convex combination of perfect matchings is computed, as above. This decomposition is then used to select a single perfect matching $\bar{M}$ from $\{M_1^{-f}, \cdots, M_k^{-f}\}$ *randomly*, where M_i^{-f} is chosen with probability λ_i^{-f} for all $i \in [k]$. Finally, the augmenting set $M \subseteq \bar{M}$ is chosen to contain all edges of $\bar{M}$ connecting *distinct connected components* of X. The edges of M are added to X and the algorithm proceeds to the next iteration. The algorithm terminates when X is a feasible solution. A summary of the algorithm is presented as Algorithm 1.

The correctness of the algorithm is shown by exploiting the following classical result for matching-covered graphs.

Theorem 11 ([42, Thm. 4.1.1., p. 122]). *A connected bipartite graph $H = (U \cup W, E)$ with $|U \cup W| \geq 4$ is matching-covered if and only if for any $u \in U$ and $w \in W$ the graph $H - \{u\} - \{w\}$ has a perfect matching.*

Concretely, Theorem 11 serves as a main ingredient for the proof of the next lemma, which states a useful structural property of intermediate solutions in the algorithm.

Lemma 12. *Let X be a non-empty set of edges already selected in an arbitrary iteration of Algorithm 1. Then, the subgraph $G[X]$ induced by X is matching-covered.*

Proof. As X is assumed to be non-empty, X contains at least one perfect matching of G. Thus, $G[X]$ spans G, i.e. $G[X] = (R \dot\cup T, X)$, and does not have isolated nodes. Let $S \subseteq R \dot\cup T$ be the nodes of some connected component of $G[X]$. It suffices to prove that the corresponding connected component $(S, X[S])$ with $X[S] := \{e \in X \mid e = \{s_1, s_2\}, \text{ for some } s_1, s_2 \in S\}$ is matching-covered.
For $|S| = 2$, $X[S]$ contains exactly one edge that belongs to a perfect matching in X. Thus, the claim is proved.
Next, assume that $|S| > 2$. To prove that $(S, X[S])$ is matching-covered, an induction on the number of iterations in Algorithm 1 is performed. In the first iteration, a perfect matching is added to X in Step 6, thus the claim holds in that case.
Now, let $X' \subseteq X$ be the set of edges selected until the beginning of the iteration preceding the current iteration. By the induction hypothesis, it can be assumed that every connected component of $(R \dot\cup T, X')$ is matching-covered. To prove the claim, it must be shown that every edge $e \in X[S]$ is contained in some perfect matching of S. If $e \in X'$, the claim holds by the inductive hypothesis, and due to $X' \subseteq X$. If $e \notin X'$, then $e \in \bar{M}$, where $\bar{M}$ is the matching selected in Step 5 in the current iteration. This means that e connects nodes from two distinct connected components of $(R \dot\cup T, X')$ since e was added in Step 6.
Now, pick any cycle $C \subseteq X$ in $G[X]$ containing e with a minimum number of edges from $\bar{M}$. Let $D_1, \cdots, D_l \subseteq R \dot\cup T$ be the components in $(R \dot\cup T, X')$ that have edges in C. From minimality of $|C \cap \bar{M}|$ it follows that C is a simple cycle (i.e. each node is contained in at most two of its edges) and that each component D_j, $j = 1, \cdots, l$ contains exactly two nodes incident to the cycle. This cycle can now be used to demonstrate the existence of the desired perfect matching M' as follows. First, include in M' all edges in $C \cap \bar{M}$. Then, in every component D_j for $j = 1, \cdots, l$ pick a matching covering all nodes, except the two nodes incident to the cycle C. Due to Theorem 11, such a matching

exists since each component D_j is matching-covered. The matching chosen so far covers exactly the nodes in $D_1 \cup \cdots \cup D_l$. Finally, pick any matching covering all other components of $(R \dot\cup T, X')$ that are not incident to C. This matching exists, since again, $(R \dot\cup T, X')$ is matching-covered. The result is a perfect matching in $G[X]$ containing e, which completes the proof. ■

Lemma 12 guarantees that at every stage in the algorithm, the only edges not yet covered by the current solution X are the isolated edges of $G[X]$. Now, since at an iteration where an uncovered edge f is chosen in Step 4, the set M must contain two edges distinct from f, that are incident to the endpoints of f, the edge f is guaranteed to be covered in the end of this iteration. This immediately implies that the algorithm terminates with a feasible solution after at most $|E|$ iterations. It hence remains to bound the expected cost of the solution returned by Algorithm 1.

Lemma 13. *The expected cost of the solution returned by Algorithm 1 is* $O(\log n) \cdot \mathrm{OPT}$*, where* OPT *is the optimal solution value for the* RAP *instance.*

Proof. The feasibility of the obtained solution and the bound on the running time are guaranteed by Lemma 12.
For a set $Q \subseteq E$ of edges, let $c_{\mathrm{LP}}(Q)$ denote the contribution of the edges in Q to the LP cost, i.e. $c_{\mathrm{LP}}(Q) = \sum_{e \in Q} c_e y_e$. For a node $v \in R \dot\cup T$, $\delta(v) \subseteq E$ represents the set of edges incident to v. To bound the approximation guarantee, the expectation of the ratio $c(X)/c_{\mathrm{LP}}(E)$ is bounded accordingly. Since the LP is a relaxation of the problem, $c_{\mathrm{LP}}(E) \leq \mathrm{OPT}$ holds. Thus, this ratio is a valid bound on the approximation guarantee.
To obtain the desired bound, a scheme charging every selected edge in any stage of the algorithm to one of its endpoints is developed. It is then shown that the expected cost charged to any node $v \in V$ is bounded by $O(\log n)$ times the *fractional cost* $c_{\mathrm{LP}}(\delta(v))$ associated with the node. This implies that the expected cost of all edges added by the algorithm is at most

$$O(\log n) \cdot \sum_{v \in R \dot\cup T} c_{\mathrm{LP}}(\delta(v)) \leq O(\log n) \cdot \mathrm{OPT},$$

where the last inequality follows from linearity of expectation, $c_{\mathrm{LP}}(E) \leq \mathrm{OPT}$ and $c_{\mathrm{LP}}(E) = \frac{1}{2} \sum_{v \in R \dot\cup T} c_{\mathrm{LP}}(\delta(v))$.
Next, it is described how the costs of the selected edges are charged to the nodes of the graph. Let $\bar{X} \subseteq E$ be the set of edges returned by the algorithm. Formally, with each node $v \in R \dot\cup T$, a collection of edges $\bar{X}_v \subseteq \bar{X}$ is associated

such that $\bigcup_{v \in R \dot\cup T} \bar{X}_v = \bar{X}$ holds and such that $c(\bar{X}_v)$ is bounded by $O(\log n)$ times the fractional load at v in expectation.
The sets $\bar{X}_v$ are constructed as follows. In the beginning $\bar{X}_v = \emptyset$ for all $v \in R \dot\cup T$. Let X be the set of edges selected so far by the algorithm and let $M \subseteq E \setminus X$ be the set of edges selected to be added to X in Step 6 of the current iteration. At this stage, the sets $\bar{X}_v$ might already contain some edges. Depending on the selection of M, the sets $\bar{X}_v$ now change as follows. Consider an edge $f = \{r, t\} \in M$. Recall that the algorithm only includes edges in the solution if they connect different connected components in $(R \dot\cup T, X)$. Thus, r and t lie in different connected components. Let D_r and D_t be the node sets of the connected components to which r and t belong, respectively, and assume without loss of generality that $|D_r| \leq |D_t|$. Then, f is charged to r, i.e. f is included in $\bar{X}_r$. In other words, an edge added by the algorithm in any iteration is charged to the node contained in the *smaller connected component*, with ties broken arbitrarily.
It is obvious that the latter scheme charges all edges in $\bar{X}$ to some nodes, such that $\bigcup_{v \in V} \bar{X}_v = \bar{X}$ holds in the end of the last iteration. It remains to analyze the quantity $c(\bar{X}_v)$ for a single node $v \in R \dot\cup T$. The bound on $c(\bar{X})$ will then follow from linearity of expectation and the previous discussion. To arrive at the desired bound it suffices to make the following two observations.
First, at any time, if an edge is charged to v, its expected cost is at most $c_{\mathrm{LP}}(\delta(v))$. Indeed, recap that the edges in M come from a perfect matching chosen at random from the decomposition of some fractional perfect matching x^{-f} in the graph (this x^{-f} corresponds to the edge f chosen in Step 4 in the current iteration). Let this decomposition be

$$x^{-f} = \sum_{i \in [k]} \lambda_i^{-f} \chi^{M_i^{-f}}.$$

The distribution over the integral matchings defining x^{-f} induces a distribution over the edges incident to v: each edge $e \in \delta(v)$ is contained in the perfect matching with a probability $p_e \in [0, 1]$ given by

$$p_e = \sum_{i \in [k] \,:\, e \in M_i^{-f}} \lambda_i^{-f} = x_e^{-f}.$$

Since $x_e^{-f} \leq y_e$ for all $e \in E$, the expected cost of the edge charged to v is at most $\sum_{e \in \delta(v)} c_e x_e^{-f} \leq c_{\mathrm{LP}}(\delta(v))$, proving the first property.
The second observation concerns the number of times the node v is charged in the course of the algorithm. Consider any iteration in which some edge was

charged to v, and let D_v be the nodes in the component of v in the beginning of the iteration. Since an edge is only charged to a node of the smaller component, and since charged edges always merge connected components, the size of the connected component containing v in the end of the iteration is *at least* $2|D_v|$. Since the graph only contains n nodes, this doubling can only happen at most $\log n$ times.
Therefore, $c(\bar{X}_v)$ is, in expectation, indeed at most $O(\log n) \cdot c_{LP}(\delta(v))$, which concludes the proof. ■

The proof of Lemma 13 uses a discharging argument that assigns the cost of selected to nodes in the graph, in a way that the total assigned cost to any node can be bounded by the fractional contribution of edges incident to the node. While it seems difficult to perform this type of argument a posteriori for a computed feasible solution, a cost assignment scheme is designed, that is constructed by following the progress of the algorithm.

Lemma 12 and Lemma 13 already imply the correctness of Theorem 4 for the uniform case. The generalization to the non-uniform case is explained in the following proof.

Proof of Theorem 4. It remains to show how to treat the case $\mathcal{F} \neq E$. For this, a transformation to reduce such an instance to a uniform instance by losing only a factor of 2 in the approximation guarantee is provided.
The transformation first adds to the graph one parallel edge $\bar{e}$ for every $e \notin \mathcal{F}$. Let G' be the obtained graph with edge set E'. The new set of vulnerable edges is set to $\mathcal{F}' = E'$. Solutions for the two RAP instances are in the following correspondence: A solution $X \subseteq E$ to the original instance on G can be transformed to a solution for the new instance on G' of at most double the cost by taking $X' = X \cup \{\bar{e} \mid e \in X \setminus \mathcal{F}\}$. Conversely, a solution X' for the new instance can be transformed to a solution for the original instance with the same, or better cost, by choosing $X = X' \cap E$.
Let OPT' denote the optimal solution value of the transformed uniform instance. Our $O(\log n)$-approximation algorithm for the general case proceeds by first transforming the instance to a uniform instance of RAP, as above, then invoking Algorithm 1 to obtain the set X', having expected cost at most $O(\log n)\text{OPT}' = O(\log n)\text{OPT}$ and then returning $X = X' \cap E$. ■

The section is concluded by arguing why simpler randomized rounding techniques, for instance, ones that lead to logarithmic approximation to many covering problems, are unlikely to lead to a similar result for RAP. The reason for

this is that there does not seem to be a simple way to obtain a compact set cover-type representation of RAP without losing a super-logarithmic factor in the approximation guarantee. One natural attempt could be to consider every vulnerable edge $f \in \mathcal{F}$ as an *element* that needs to be covered, and every possible perfect matching $M \subseteq E$ that does not contain f, a *covering set* that covers the edge f (and all other edges in $\mathcal{F} \setminus M$). The cost of the covering set is simply the sum of the costs of edges in the corresponding perfect matching. Unfortunately, it is easy to come up with examples in which the optimal solution value in the latter set covering model has cost $\Omega(n)$OPT. Such instances can be constructed, for example by choosing an instance, such that any feasible solution must have some nodes with very high degree, while an optimal solution has cost $O(n)$.

4 Hardness and approximability of card-RAP

In this section we restrict our attention to the special case card-RAP, i.e. to the problem of finding a robust assignment of minimum cardinality. Subsection 4.1 deals with the proof of Theorem 5 stating that card-RAP is hard to approximate up to some constant $\delta > 1$. This implies thtat card-RAP does not admit a PTAS unless P = NP. An $O(1)$-approximation algorithm for card-RAP is then presented in Subsection 4.2, thus proving Theorem 6.

4.1 Hardness of Approximation for card-RAP

As is the case for general uniform RAP, the hardness proof we present here for card-RAP is also based on a reduction from the Set Cover Problem (SCP).

By Lemma 10, any given SCP instance $\mathcal{I} := ([k], \mathcal{S})$ can be equivalently transformed, in time polynomial in the input size of $\mathcal{I}$, into a uniform RAP instance $\bar{\mathcal{I}} := (\bar{G}, \bar{E}, \bar{c})$. Note that $\bar{G} := (\bar{R} \dot{\cup} \bar{T}, \bar{E})$ is the graph constructed from $\mathcal{I}$ by applying the transformation steps **(T1)**–**(T4)** as stated on page 15 and step **(T5)** from page 18, whereas the cost vector $\bar{c} \in \mathbf{R}^{\bar{E}}_{\geq 0}$ is chosen with $\bar{c}_e = 1$, if e belongs to $\bar{E}_4 = E_4$, and $\bar{c}_e = 0$, otherwise. Recap further that the specific choice of $\bar{c}$ is crucial for deriving an equivalent reformulation of $\mathcal{I}$ in terms of a uniform RAP. Lemma 10 (b) indeed shows that all edges from $\bar{G}$ except those from E_4 can be assumed to be contained in an optimal solution to $\bar{\mathcal{I}}$ as they have zero cost. This property does not hold anymore when $\bar{\mathcal{I}}$ is replaced with its unweighted version.

In order to arrive at a card-RAP instance being equivalent to the given SCP instance $\mathcal{I}$, a further transformation step to be performed on $\bar{G}$ is needed.

(T6) For each edge $f_s = \{\bar{u}_s, u_s\} \in E_1$ subdivide the edge by adding two new nodes z_1^s and z_2^s.

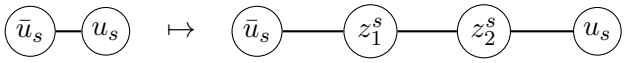

Figure 3: Replacing an edge $f_s = \{\bar{u}_s, u_s\} \in E_1$ by a path of length three.

Transformation step **(T6)** is illustrated in Figure 3.

In the following, let $\tilde{G} = (\tilde{R} \dot{\cup} \tilde{T}, \tilde{E})$, with $\tilde{E} = \bigcup_{j=1}^{6} \tilde{E}_j$, denote the graph obtained by applying transformation step **(T6)** to $\bar{G}$, and let $\tilde{\mathcal{I}}$ be the corresponding uniform card-RAP instance defined on $\tilde{G}$, i.e. $\tilde{\mathcal{I}} = (\tilde{G}, \tilde{E}, \mathbf{1}_{\tilde{E}})$ where $\mathbf{1}_E$ is the all one vector in $\mathbf{R}^{\tilde{E}}$. Observe that

- $\tilde{R} = \bar{R} \cup \{z_2^s \mid s \in [k]\}$, $\tilde{T} = \bar{T} \cup \{z_1^s \mid s \in [k]\}$,
- $\tilde{E}_1 := \big\{ \{\bar{u}_s, z_1^s\}, \{z_1^s, z_2^s\}, \{z_2^s, u_s\} \mid s \in [k] \big\}$,
- $\tilde{E}_3 = \bar{E}_3$, $\tilde{E}_5 = \bar{E}_5$,
- $\tilde{E}_2 = \bar{E}_2 = E_2$, $\tilde{E}_4 = \bar{E}_4 = E_4$, and $\tilde{E}_6 = \bar{E}_6 = E_6$

hold. In order to formulate an analogue of Lemma 9 the following notion of *efficiency* for feasible solutions to $\tilde{\mathcal{I}}$ is useful.

Concretely, a feasible solution $\tilde{X} \subseteq \tilde{E}$ to $\tilde{\mathcal{I}}$ is called *efficient*, if $|\tilde{X} \cap \tilde{E}_2| = |\tilde{X} \cap \tilde{E}_6| = k$. To justify this definition, consider any non-efficient feasible solution X. On the one hand, feasibility of X implies that $|X \cap \tilde{E}_2| \geq k$, as every node u_i, $i \in [k]$, must have at least one incident edge in $\tilde{E}_2$, that is contained in X. Similarly, $|X \cap \tilde{E}_6| \geq k$ must hold. On the other hand, it is next argued that X can be transformed in polynomial time into a new feasible solution $\tilde{X}$ that is efficient.

By a similar argument as used in the proof of Lemma 9, it can be shown that $\mathcal{C}_X = \{S \in \mathcal{S} \mid \{\bar{v}_S, \tilde{v}_S\} \in X\}$ is a feasible cover to the given SCP instance $\mathcal{I}$. Now, for each $i \in [k]$, choose an arbitrary set $S_{j_i} \in \mathcal{C}_X$ with $i \in S_{j_i}$. Then,

$$\tilde{X} := \big(X \setminus (\tilde{E}_2 \cup \tilde{E}_6) \big) \cup \big\{ \{u_i, v_{S_{j_i}}\}, \{\bar{u}_i, w_{S_{j_i}}\} \mid i \in [k] \big\}$$

is a feasible to $\tilde{\mathcal{I}}$ with $|\tilde{X}| \leq |X|$. In addition, $\tilde{X}$ is efficient and can be constructed in polynomial time.

Lemma 14. *Let $\mathcal{I} = ([k], \mathcal{S})$ be any* SCP *instance, and let $\tilde{\mathcal{I}}$ be the corresponding uniform* card-RAP *instance on the graph $\bar{G} := (\tilde{R} \dot\cup \tilde{T}, \tilde{E})$ that is obtained by applying the transformation steps* **(T1)**–**(T6)** *on $\mathcal{I}$. Moreover, let $q := |\tilde{E}_1 \cup \tilde{E}_3 \cup \tilde{E}_5|$.*

(a) $\tilde{\mathcal{I}}$ can be constructed from $\mathcal{I}$, in polynomial time.

(b) Let $\tilde{X} \subseteq \tilde{E}$ be an efficient feasible solution to $\tilde{\mathcal{I}}$. Then there exists a feasible cover for $\mathcal{I}$, that is of size $|\tilde{X}| - q - 2k$.

(c) If $\mathcal{C}$ is a feasible cover for $\mathcal{I}$, then there exists a feasible solution $\tilde{X}$ to $\tilde{\mathcal{I}}$ with $|\tilde{X}| = |\mathcal{C}| + q + 2k$.

Proof.

(a) This statement follows from Lemma 10 (a) and the fact that performing transformation step **(T6)** on $\bar{G}$ can also be done in polynomial time.

(b) Let $\tilde{X} \subseteq \tilde{E}$ be an efficient, feasible solution. Since every edge from $\tilde{G}$ is vulnerable and since every edge in $\tilde{E}_1 \cup \tilde{E}_3 \cup \tilde{E}_5$ is incident to some node of $\tilde{G}$ with degree two, it follows that $\tilde{E}_1 \cup \tilde{E}_3 \cup \tilde{E}_5 \subseteq \tilde{X}$ must hold. Furthermore, $|(\tilde{E}_2 \cup \tilde{E}_4) \cap \tilde{X}| = 2k$ implies that $\tilde{X}$ contains $|\tilde{X}| - q - 2k$ edges from $\tilde{E}_4$. Thus, $\mathcal{C}_{\tilde{X}} = \left\{S \in \mathcal{S} \mid \{\bar{v}_S, \tilde{v}_S\} \in \tilde{X}\right\}$ is a feasible cover for $\mathcal{I}$ with $\mathcal{C}_{\tilde{X}} = |\tilde{X}| - q - 2k$.

(c) Let $\mathcal{C}$ be a feasible cover for $\mathcal{I}$. For each $i \in [k]$, let further denote by j_i the index of some set of $\mathcal{S}$ that covers i in $\mathcal{C}$. Then,

$$\tilde{X} := \tilde{E}_1 \cup \tilde{E}_3 \cup \tilde{E}_5 \\ \cup \{\{u_i, v_{S_{j_i}}\}, \{\bar{u}_i, w_{S_{j_i}}\} \mid i \in [k]\} \cup \{\{\bar{v}_{S_j}, \tilde{v}_{S_j}\} \mid S_j \in \mathcal{C}\}$$

is a feasible solution to $\tilde{\mathcal{I}}$, that additionally satisfies $|\tilde{X}| = |\mathcal{C}| + q + 2k$.

■

Now, Theorem 5 can be proved with the help of Lemma 14.

Proof of Theorem 5. It is well known that the *Vertex Cover Problem in subcubic Graphs* (VC 3) on an input graph $H = (V, E_H)$ with $|E_H| = k$ can be equivalently restated as an instance of the Set Cover Problem, where the ground set is $E_H \cong [k]$ and each $S \in \mathcal{S}$ corresponds to a cut set $\delta(v)$, for some $v \in V$. Note that $|S| \leq 3$, for all $S \in \mathcal{S}$, and $|\{S \in \mathcal{S} \mid s \in S\}| = 2$ for all $s \in [k]$. Alimonti and Kann proved in [8] that there exists a constant $\delta > 1$ such that

VC 3 does not admit a polynomial δ-approximation algorithm unless P=NP.
Now, let a VC 3 instance be given in terms of a SCP instance $([k], \mathcal{S})$. Let $\tilde{G}$ be the graph obtained form $([k], \mathcal{S})$ by applying the transformation steps **(T1)**–**(T6)**, let $\tilde{\mathcal{I}}$ be the uniform card-RAP instance induced by $\tilde{G}$.
Observe that $\tilde{G}$ has $O(k)$ edges and that any feasible solution to VC 3 contains at least $k/3 = \Omega(k)$ sets. From Lemma 14, it follows that for any constant $\alpha > 1$ there exists a constant $\beta = \beta(\alpha) > 1$ such that any polynomial β-approximation algorithm for the uniform card-RAP can be used to construct a polynomial α-approximation algorithm for VC 3. Since no δ-approximation exists for VC 3, unless P=NP, there is also no $\alpha(\delta)$-approximation for the uniform card-RAP unless P=NP. ■

4.2 $O(1)$-Approximation for card-RAP

In this subsection, a constant factor approximation algorithm for card-RAP is developed. The algorithm makes use of an ear decomposition of the underlying bipartite graph. At the end of this subsection it is also shown that the algorithm presented can be easily extended to unbalanced graphs.
Note that similar approaches using ear decompositions have successfully been used to approximate various combinatorial optimization problems. Among others, these works include algorithms for the minimum edge connected subgraph problem by Cheriyan, Sebő and Szegeti [18] as well as the traveling salesman problem by Sebő and Vygen [48].

Definition 15 (Ear Decomposition of Bipartite Graphs (e.g., p. 123 in [42]))**.** Let H be a bipartite graph, and let H' be a subgraph of H. An *odd ear of H with respect to H'* is a path P in H with an odd number of edges and such that P and H' have exactly two nodes in common. Those two nodes form the end points of P, and belong to different parts of the bipartition.
A *bipartite ear decomposition of a bipartite graph H* is a sequence $P_0, P_1, \ldots, P_q$ of paths in H, such that:

(i) $P_0 = (\{v_1, v_2\}, \{\{v_1, v_2\}\})$ is a graph composed of a single edge,

(ii) $H = P_0 + \cdots + P_q$,

(iii) for every $j = 1, \ldots, q$, the path P_j is an odd ear with respect to $H_{j-1} := P_0 + \cdots + P_{j-1}$.

The next theorem provides a well-known connection between matching-covered bipartite graphs and bipartite ear decompositions.

Theorem 16 ([42, Thm. 4.1.6])**.** *A bipartite graph is matching-covered if and only if it has a bipartite ear decomposition.*

Now consider an instance $\mathcal{I} = (G, \mathcal{F})$ of card-RAP with uncertainty set $\mathcal{F} \subseteq E[G]$. As always, we assume that $\mathcal{I}$ is feasible. In order to apply the theory of matching-covered graphs the algorithm first removes from G all so-called *dispensable edges*, i.e. all edges not appearing in any perfect matching of G. This way, a graph is obtained that is, by definition, matching-covered. Note that the removal can be implemented in polynomial time using any efficient algorithm for finding bipartite matchings. We also assume that the removal of the dispensable edges results in a connected graph (otherwise we treat each connected component separately). Moreover, recap that dispensable edges can be always removed from any feasible solution without breaking feasibility. In the following, the new graph resulting from deleting all dispensible edges is also called G to facilitate readability.

Now let $G = P_0 + \cdots + P_q$ be any ear decomposition of G (which is by no means unique) with some arbitrary chosen initial edge P_0. An ear P_j of the decomposition is called *trivial* if it is not P_0 and if it consists of a single edge, only. The next lemma shows that a feasible solution to $\mathcal{I}$ can be obtained from the ear decomposition of G by skipping trivial ears.

Lemma 17. *Let $J := \{j \in [q] \mid P_j$ is a trivial ear$\}$, and let G' be the subgraph in G resulting from G by removing all trivial ears, i.e. $G' = P_0 + \sum_{i \in [q] \setminus J} P_i$. Then, $X := E[G']$ is a feasible solution to the* card-RAP *instance on the original graph $G = (R \dot\cup T, E)$. Furthermore, $|X| \leq 3|T|$ holds.*

Proof. The subgraph G' is obtained from G by deleting trivial ears only, i.e. G' is spanning in G. By Proposition 2, it suffices to show that G' is matching-covered and that there are no isolated edges. The fact that G' is matching-covered follows from the fact that it has, by definition, an ear decomposition $G' = P_0 + \sum_{i \in [q] \setminus J} P_i$ and from Theorem 16. Since G has a feasible solution to card-RAP, there are no isolated vulnerable edges in G, and hence also not in G'. This proves that $X = E'$ is feasible.

It remains to bound the number of edges in $X = E[G']$. Let $G' = Q_0 + \ldots + Q_p$ be the ear decomposition of G' consisting of the non-trivial ears of G (appearing in the same order) with $p := q - |J|$. Furthermore, define $l_j := |T \cap (V[Q_j] \setminus (V[Q_0] \cup \cdots \cup V[Q_{j-1}]))|$ as the number of internal task nodes in the ear Q_j for

$j = 1, \ldots, p$. Since the ears Q_j are non-trivial we have $l_j \geq 1$, $p \leq |T| - 1$ and $\sum_{j=0}^{p} l_j = |T|$ implying

$$|X| \leq 1 + \sum_{j=1}^{p} 2(l_j + 1) = 1 + 2|T| + p \leq 3|T|.$$

■

Lemma 17 allows to arrive at an approximation algorithm summarized as Algorithm 2. The analysis is deferred to the proof of Theorem 6.

Algorithm 2 : $O(1)$-Approximation for card-RAP

Require: $G = (R \dot{\cup} T, E)$ and $\mathcal{F} \subseteq E$.
Ensure: a feasible solution X to the card-RAP instance on G and $\mathcal{F}$
1: $X \leftarrow \emptyset$
2: Remove all dispensable edges form G
3: Compute an ear decomposition $G = P_0 + \ldots + P_q$
4: $X \leftarrow P_0 \cup \bigcup\{E[P_j] \mid P_j \text{ is not trivial, } j = 1, \ldots, q\}$
5: **return** X

Proof of Theorem 6. According to de Carvalho and Cheriyan [22], an ear decomposition of matching-covered graphs can be computed in polynomial time. Furthermore, all other computations can also be implemented efficiently, such that the running time of Algorithm 2 is polynomial.
From Lemma 17, it follows that the set X returned by Algorithm 2 is feasible. For $\mathcal{F} = E$, any feasible solution must have at least two edges incident to any node from the set T. Hence, $\text{OPT} \geq 2|T|$. Since $|X| \leq 3|T|$, the approximation guarantee is indeed 1.5. If $\mathcal{F} \subsetneq E$, then G can contain a perfect matching not including any edge from $\mathcal{F}$, hence $\text{OPT} \geq |T|$. Using $\text{OPT} \geq |T|$ yields an approximation factor of 3. ■

In the example below, we provide a family of graphs, for which Algorithm 2 produces a sequence of solutions whose objective function values can be arbitrarily close to $1.5\,\text{OPT}$ in the uniform case.

Example 18. For $k \geq 3$, let G_k be a bipartite graph with node set $\{0, 1, \ldots, 2k+1\}$. The edges of G_k are defined as follows. G_k contains the edge $\{0, 1\}$ as well as the paths P_i, $i \in \{2z \mid z = 1, \ldots, k\}$, from 0 to 1 through nodes i and $i + 1$.

Additionally, G_k contains the cycles C_j, $j \in \{2z \mid z = 2, \ldots, k-1\}$ with node set $\{j, j+1, j+2, j+3\}$. Figure 4 shows the graph G_k, for $k = 3$.

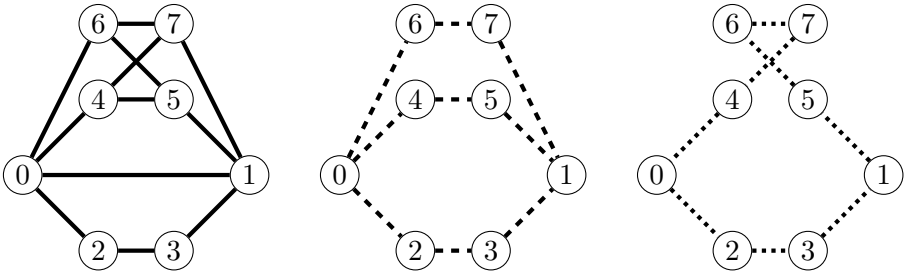

Figure 4: Graph G_3 (left), a bad solution found by Algorithm 2 (center), an optimal solution for $\mathcal{F} = E$ (right).

In case $\mathcal{F} = E$, an optimal solution of the corresponding card-RAP instance is given by a Hamiltonian cycle of size $2k + 2$, while a "worst case" solution that can be found by Algorithm 2 is given by the union of all paths P_i, i.e. by $\bigcup_{i \in \{2z \mid z=1,\ldots,k\}} P_i$. This yields

$$\frac{\text{ALG}}{\text{OPT}} = \frac{3k}{2(k+1)} \xrightarrow[\infty]{k} \frac{3}{2}.$$

It is worth to remark that Algorithm 2 presented above can also be applied to unweighted RAP instances defined on unbalanced bipartite graphs, leading to the same approximation guarantees. The remainder of this section explains the extension of the algorithm.

Consider a bipartite graph $G_u = (R_u \dot\cup T_u, E_u)$ with $|T_u| < |R_u|$. As an initial pre-processing step, the graph G_u is transformed into a balanced bipartite graph G_b (cf. Proposition 1). Therefore a set D of $|R_u| - |T_u|$ dummy task nodes is introduced. Further, each such newly introduced node is then connected with all nodes from R_u, yielding a new subclass E_D of edges. The desired balanced bipartite graph is then given by $G_b := (R_u \dot\cup (T_u \cup D), E_u \cup E_D)$ and defines a card-RAP instance where the uncertainty set is chosen to be $\mathcal{F}_b = \mathcal{F}_u \cup E_D$. As before all dispensable edges has to be removed from G_b.

In Lemma 17 the set X is now defined as $X := E[G'] \setminus E_D$. $E[G']$ is feasible for the instance on G_b and X remains feasible for G_u because the edges E_D are incident only to dummy task nodes in D and not to those in T_u. Thus, $X \setminus \{f\}$ still contains a matching covering the whole set T_u. The second part

of the proof of Lemma 17 holds if T is replaced by T_u. The same is true for the proof of Theorem 6. Hence, the approximation guarantee remains the same for an unbalanced bipartite graph G_u.

5 Hardness of card-RAP with two vulnerable edges

In this section the proof of Theorem 7, stating that card-RAP with only two vulnerable edges is NP-hard, is presented. The proof requires a complexity analysis of the following optimization problem.

Problem (Shortest Nice Path Problem (SNPP))**.**

- Input: Tuple (H, s, t), where $H = (U \dot\cup W, E_H)$ is a balanced bipartite graph, and $s \in U$, $t \in W$.
- Output: A nice[5] shortest s-t-path in H, i.e. an s-t-path P with the smallest possible number of edges such that there is a a matching in H covering all nodes in H that are not covered by P.

The proof of Theorem 7 is based on the fact that SNPP is NP-hard. This result is later shown in Lemma 19. We first show how Theorem 7 follows from the latter result.

Proof of Theorem 7. It suffices to show that SNPP can be reduced, in polynomial time, to card-RAP with only two vulnerable edges. Then, the NP-hardness of SNPP shown in Lemma 19 completes this proof.
Consider any instance $\mathcal{I} := (H, s, t)$ of SNPP with graph $H := (U \dot\cup W, E_H)$ and with terminals $s \in U$ and $t \in W$. Further, set $n := |U| + |W|$. To derive an equivalent card-RAP instance, two new nodes x, y are introduced as well as the edges $f_1 := \{s, x\}$, $f_2 := \{x, y\}$ and $g := \{y, t\}$. For an illustration, see Figure 5 This leads to a new balanced bipartite graph $G = (T \dot\cup R, E)$ with $T := U \cup \{y\}$, $R := W \cup \{x\}$ and $E := E_H \cup \{f_1, f_2, g\}$.

Now consider the card-RAP instance $\mathcal{I}'$ defined on G and with uncertainty set $\mathcal{F} = \{f_1, f_2\}$. $\mathcal{I}'$ can clearly be constructed in polynomial time from $\mathcal{I}$.
Let P be any nice s-t-path P in H. As $s \in U$ and $t \in W$, P consists of an even number of nodes and an odd number of edges. The definition of nice

[5] Recap from the theory of matching-covered graphs that a subgraph H' of a graph H containing a perfect matching is called *nice* if $(V[H] \setminus V[H'])$ is perfectly matchable in H.

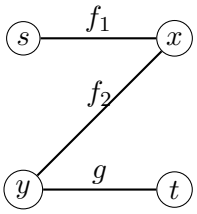

Figure 5: Additional nodes and edges in the graph G.

paths implies that there must exist a matching M_P in H covering all nodes $(U \dot\cup W) \setminus V[P]$, i.e. all nodes from H not belonging to P.
In the graph G, the path P and the new edges f_1, f_2, g form the cycle $C := P + (\{x, y\}, \{f_1, f_2, g\})$ with even length. Let $E[C]$ be the set of edges contained in C.
It is next argued that the union $M_P \cup E[C]$ is a feasible solution to the card-RAP instance $\mathcal{I}'$. Indeed, as the cycle C is even, its edge set $E[C]$ is the union of two perfect matchings M_1 and M_2. This in particulary means that both M_1 and M_2 cover all nodes from C. Since the vulnerable edges f_1, f_2 are adjacent on C, one of these matchings contains f_1, while the other one contains f_2. W.l.o.g. assume that $f_1 \notin M_1$ and $f_2 \notin M_2$. Now recap that M_P covers all nodes from G not contained in P. Then, for each f_i, $i = 1, 2$, $M_i \cup M_P$ is a perfect matching in G not containing f_i. This shows that $M_P \cup E[C]$ is feasible to $\mathcal{I}'$.
Summing up, a nice s-t-path P in H with L nodes (and $L - 1$ edges) can be used to construct a feasible solution to $\mathcal{I}'$ with $L + 2 + \frac{n-L}{2} = \frac{n}{2} + \frac{L}{2} + 2$ edges.

On the other hand, it is claimed that the following statements are true.

(i) Any solution X feasible to the card-RAP instance $\mathcal{I}'$ must contain a cycle C including f_1 and f_2.

(ii) If X is *optimal* to $\mathcal{I}'$ then it must addionally hold that all edges in X not belonging to C form a matching covering all remaining nodes.

To show the correctness of both statements, observe first that f_1 and f_2 are contained in any feasible set X as they are the only edges in G being incident with node x. Now, let $M_1, M_2 \subseteq X$ be two perfect matchings in G with M_i not containing f_i, for $i = 1, 2$. Since x has degree two in G it holds that $f_2 \in M_1$ and $f_1 \in M_2$. Hence, $M_1 \cup M_2 \subseteq X$ must contain a cycle involving both f_1 and f_2, implying that statement (i) is true.
If X is optimal, it further follows that $X = M_1 \cup M_2$ holds. Recap that a union of perfect matchings is, in general, a union of even cycles with a matching. Thus,

it still needs to be verified that X cannot contain other cycles than C. For this, assume there is a second cycle $C' \neq C$ in X. Since both f_1 and f_2 must be contained in the same cycle, the existence of C' implies that C' contains neither f_1 nor f_2. As C' is also even, it contains two different matchings covering its nodes. Then, one of those matching can be deleted from X, yielding a proper subset of X that is still feasible to $\mathcal{I}'$. This contradicts optimality of X and shows the correctness of statement (ii).

Finally, since y has also a degree of two in G, edge g has to be part of the cycle C containing f_1 and f_2. It follows that $C - \{x, y\}$ is a nice s-t-path in H. Consequently, if for some integer L, the cardinality of an optimal solution to the card-RAP instance is $\frac{n}{2} + \frac{L}{2} + 2$, the cycle contained in this solution has $L + 2$ nodes, and thus it provides a nice s-t-path with L nodes.

To conclude, H admits a nice s-t-path with L nodes if and only if the optimal solution of the card-RAP instance on G with f_1, f_2 being vulnerable has cardinality of $\frac{n}{2} + \frac{L}{2} + 2$. This shows that SNPP is reducible to card-RAP with two vulnerable edges, in polynoimal time. ■

The remainder of this section addresses the NP-hardness of SNPP stated in the next lemma.

Lemma 19. SNPP *is* NP*-hard.*

The proof of Lemma 19 resorts to a reduction from a special variant of the following well-known NP-complete decision problem.

Problem (Path with Forbidden Pairs Problem (PFP), [31, GT54])**.**

- Input: Tuple $(D, s, t, \{\{u_i, v_i\} \mid i \in \{1, \ldots, k\}\})$, where $D = (V, A)$ is a directed graph, s, t are two terminals in V, and $\{u_i, v_i\}$ are k pairs of nodes in V (with all $2k$ nodes being distinct).
- Question: Does D contain an s-t-path P with the property that, for every $i \in [k]$, at most one node from $\{u_i, v_i\}$ is contained in P?

To show NP-hardness of SNPP the following restricted version of PFP is considered.

Problem (Restricted Path with Forbidden Pairs Problem (RPFP))**.**

- Input: Tuple $(H, s, t, \{\{u_i, v_i\} \mid i \in \{1, \ldots, k\}\})$, where

(P1) $H = (U \dot\cup W, E_H)$ is undirected, bipartite and balanced.

(P2) $s \in U$ and $t \in W$.

(P3) k is even.

(P4) $\{u_i, v_i\} \subseteq U$ or $\{u_i, v_i\} \subseteq W$ for all $i \in [k]$.

(P5) $\big|\{i \in \{1, \dots, k\} \mid u_i, v_i \in U\}\big| = \big|\{i \in \{1, \dots, k\} \mid u_i, v_i \in W\}\big| = \frac{k}{2}$.

- Question: Does H contain an s-t-path P with the property that, for every $i \in [k]$, at most one node from $\{u_i, v_i\}$ is contained in P?

It turns out that the special version RPFP remains NP-complete.

Lemma 20. RPFP *is* NP*-complete.*

Proof. To prove this statement, the proof of the NP-completeness result for PFP as given by Gabow et al. [29, Lemma 2] can be adapted because of which it is only sketched here.[6]
The main idea is to use a reduction from the NP-complete *3-Satisfyability Problem* 3-SAT. This is achieved by modeling pairs of a literal and its negation as forbidden pairs. For each clause appearing in the Boolean formula that specifies the 3-SAT instance, a layer of nodes is further introduced. Each two neighboring layers are then connected via a complete bipartite graph. A path avoiding the forbidden pairs yields a truth assignment and vice versa.
The main difference to the proof given in [29, Lemma 2] is that, for each clause in the Boolean formula, two identical layers of nodes needs to be introduced in order to satisfy the properties (P4) and (P5) of RPFP. ■

The section is concluded by presenting the proof on the NP-hardness of SNPP.

Proof of Lemma 19. As already mentioned, the proof is based on a polynomial reduction from RPFP to SNPP that is shown next.
Let $H = (U \dot\cup W, E_H)$ be a balanced bipartite graph, $s \in U, t \in W$, and let $\{(u_1, v_1), \dots, (u_k, v_k)\}$ be a collection of forbidden pairs, all together comprising an instance $\mathcal{I}$ of RPFP. To obtain a corresponding SNPP instance $\mathcal{I}'$ the following five steps are performed on H, and illustrated in Figure 6.

(S1) Set $L := |U \dot\cup W| + 1$ (that is an odd number as $|U \dot\cup W|$ is even).

[6] In [29], PFP is called *Impossible Pairs Constrained Path Problem.*

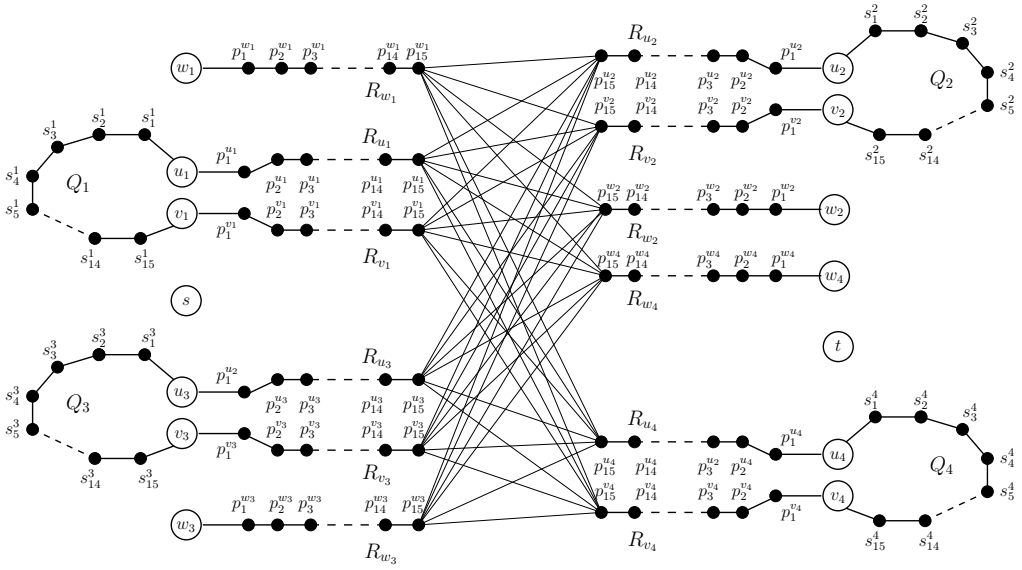

Figure 6: For a given RPFP instance $(H, s, t, \{\{u_i, v_i\} \mid i \in [4]\})$ on a balanced bipartite graph $H := (U \dot\cup W, E_H)$ with $U := \{s, u_1, v_1, u_3, v_3, w_1, w_3\}$ and $W := \{t, u_2, v_2, u_4, v_4, w_2, w_4\}$, the figure shows the graph H' derived from H by applying the transformation steps (S1)–(S5). Larger, unfilled nodes represent the original nodes from H, while the new nodes introduced through the transformation steps appear as small black nodes. To improve the readability of the figure, the edge set E_H from the original graph H is omitted. Note that $L = 15$, and recap that $s' = s$ and $t' = t$ hold.

(S2) For every $i \in \{1, \dots, k\}$, introduce new nodes s^i_j, $j = 1, \dots, L$, as well as the new path $Q_i := (u_i, s^i_1, s^i_2, \dots, s^i_L, v_i)$ connecting the two nodes of the corresponding forbidden pair $\{u_i, v_i\}$ through the new nodes s^i_j, $j = 1, \dots, L$, and add both the new nodes and the edges of Q_i to H.

(S3) For every node $w \in (U \dot\cup W) \setminus \{s, t\}$, introduce a new path $R_w := (w, p^w_1, p^w_2, \dots, p^w_L)$, and add the new nodes p^w_q, $q = 1, \dots, L$ and the edges of R_w to H.

(S4) Define $T := \{p^w_L \mid w \in (U \dot\cup w) \setminus \{s, t\}\}$, and add to H the edge set

$$E_T := \big\{ \{p^{w_1}_L, p^{w_2}_L\} \big\} \mid w_1 \in U \setminus \{s\}, w_2 \in W \setminus \{t\}\}.$$

This yields a complete, balanced, bipartite subgraph $H_T := (T, E_T)$.

(S5) Set $s' = s$ and $t' = t$.

Let $H' := (U' \dot\cup W', E_{H'})$ be the graph after performing steps (S1) to (S5). The bipartiteness of H' follows from property (P4) of RPFP, i.e. from the fact that, for each forbidden pair (u_i, v_i), both u_i and v_i belong to the same node partition of H. By property (P5) of RPFP, the number of forbidden pairs contained in each bipartition is the same. Thus, H' is also balanced. Moreover, H' contains H as a subgraph, and it may be assumed that $U \subseteq U'$ and $W \subseteq W'$, i.e. $s' \in U'$ and $t' \in W'$. Therefore, $\mathcal{I}' := (H', s', t')$ is an instance of SNPP constructed from $\mathcal{I}$ in polynomial time.
Note further that each sub-path $(s_1^i, s_2^i, \ldots, s_L^i)$ introduced in step (S2) contains an odd number of nodes. Thus, the nodes sets of the form $\{s_1^i, s_2^i, \ldots, s_L^i\}$ cannot be perfectly matched among themselves in H'. The same is true for the sub-paths $(p_1^w, p_2^w, \ldots, p_L^w)$ introduced in step (S3). Observe also that, for each $w \in T$, the nodes p_L^w and w are in different parts of the bipartition, which is due to the fact that L is odd.

Now consider the following claim.

(C) SNPP instance $\mathcal{I}'$ contains an s-t-path of length $\ell < L$ if and only if the RPFP instance $\mathcal{I}$ is feasible.

Together with Lemma 20, the correctness of claim (C) implies that finding a shortest nice path is NP-hard, thus completing the proof of this lemma.

"only if" part of the claim (C). Let P be a nice s-t-path of length $\ell < L$ in H'. Since P has at most $L - 1$ edges, it can not contain an edge from one of the paths Q_i, $i \in [k]$, and R_w, $w \in (U \dot\cup W) \setminus \{s, t\}$, as otherwise all edges from Q_i or R_w must be contained in P. This, however, implies that the length of P is greater than L as all paths Q_i and R_w consist of $L + 1$ edges (see steps (S2) and (S3)), leading to a contradiction.
Furthermore, since P contains no edge from a path R_w, it is also disjoint from the subgraph H_T introduced in step (S4). It follows that P is completely contained the subgraph H of H', i.e. P is already a path in H.
Next, it is shown that P is feasible to $\mathcal{I}$. This is a achieved by proving that, for every $i \in [k]$, at least one of the nodes u_i, v_i is not incident to P. For this, let $\{u_i, v_i\}$ be any forbidden pair and let $V[P]$ be the node set of P. Recall that any node from $V[P]$ is a node in H. Since P is a nice path in H', there exists a matching M in H' covering all nodes from H' that are not contained in $V[P]$. As the inner nodes $s_1^i, s_2^i, \ldots, s_L^i$ of any path Q_i does not belong to P, it follows that they are covered by M. Recap moreover that L is odd. This,

in particulary, means that either $\{s_1^i, u_i\} \in M$ or $\{s_L^i, v_i\} \in M$ holds. Thus, either u_i or v_i is incident to P.

"if" part of the claim (C). Let P be a feasible solution to the RPFP instance $\mathcal{I}$, i.e. P is a path in H. To show that P is a nice path in H' it must be proved that there is a matching M in H' that covers all nodes except those from $V[P]$. For this we assume w.l.o.g. that at most the nodes u_i are part of the path P and define the following index sets.

$$J := \{i \in [k] \mid u_i \in V[P]\}, \quad K := [k] \setminus J.$$

Note, that the sets J and K form a partition of $[k]$.
The matching M is constructed as follows. Set $M := \emptyset$, and consider first the nodes on the paths Q_i, $i \in [k]$. To cover all nodes of Q_i not belonging to $V[P]$, a suitable set of alternating edges (i.e. every second edge) from the path Q_i is chosen. Concretely, if $i \in J$, add to M the set of alternating edges covering $s_1^i, \ldots, s_L^i, v_i$. Otherwise, i.e. if $i \in K$, add to M the set of alternating edges that match $u_i, s_1^i, \ldots, s_L^i$. Observe that M already covers either u_i or v_i, for each forbidden pair $\{u_i, v_i\}$. By properties (P4) and (P5) of RPFP it further follows that the number of nodes in U that are covered by M and that belong to a forbidden pair is identical to the number of nodes in W covered by M and belonging to a forbidden pair.
Secondly, consider all nodes in subgraph H that do not belong to $V[P]$ and that are not covered by M, so far. Let $\tilde{U} \cup \tilde{W}$ denote these nodes where $\tilde{U} \subseteq U$ and $\tilde{W} \subseteq W$. Then, $|\tilde{U}| = |\tilde{W}|$ holds, which follows from the fact that

$$|\tilde{U}| = \frac{n}{2} - q - \frac{k}{2} = |\tilde{W}|,$$

where

- $\frac{n}{2} = |U| = |W|$ is the number of nodes on each side of H,
- $q = \frac{1}{2}|V[P]|$ (Note that $|V[P]|$ is even as $s \in U$ and $t \in W$ and P is an s-t-path in H. More precisely, even $|V[P] \cap U| = |V[P] \cap W|$ holds.),
- and $\frac{k}{2}$ the number of nodes in W (as well as in U) that are covered by M and that belong to a forbidden pair.

To extend M to a matching in H' also covering $\tilde{U} \cup \tilde{W}$, the following edges are added to M. For each $w \in \tilde{U} \cup \tilde{W}$, choose the set of alternating edges from path R_w that cover $w, p_1^w, p_2^w, \ldots, p_L^w$.
Now observe that the only nodes in H' not belonging to P and being still unmatched by M are all but the first nodes $\{p_1^w, p_2^w, \ldots, p_L^w\}$ of a path R_w that

is associated with a node $w \in (U \cup W) \setminus (\tilde{U} \cup \tilde{W} \cup \{s,t\})$, i.e. with a node w in H that is either contained in $V[P]$ or w is a node that is part of a forbidden pair and that is covered by an edge added to M in the first step. These nodes can be covered by first adding, for each $w \in (U \cup W) \setminus (\tilde{U} \cup \tilde{W} \cup \{s,t\})$, the edges $\{p_1^w, p_2^w\}$, $\{p_3^w, p_4^w\}, \ldots, \{p_{L-2}^w, p_{L-1}^w\}$ from R_w to M.
This still leaves all end nodes p_L^w of the paths R_w with a node $w \in (U \cup W) \setminus (\tilde{U} \cup \tilde{W} \cup \{s,t\})$ to be unmatched by M. Let

$$T' := \{p_L^w \mid w \in (U \cup W) \setminus (\tilde{U} \cup \tilde{W} \cup \{s,t\})\} \subseteq T$$

be the set of all these end nodes. It remains to extend M to a matching also covering T'. Recap that T' is, as a subset of T, a part of the balanced bipartite subgraph $H_T = (T, E_T)$ constructed in step (S4). Furthermore, it holds that $|T' \cap U| = \frac{1}{2}|T'| = |T' \cap W|$. Therefore, H_T contains a matching only covering nodes from T'. After adding one such matching on T' to M, the set M becomes a matching in H' that covers all nodes of H' not contained in P. This shows that P is a nice s-t-path in H'. As P is completely contained in the subgraph H, its node set $V[P]$ can only consists of at most $|U \dot\cup W| < |U \dot\cup W| + 1 =: L$ nodes, proving that the length ℓ of P is strictly less than L. ■

6 Conclusion and Future Work

This paper studies a novel practically relevant robust variant of the assignment problem (RAP). Tight connections between RAP and classical notions in matching theory, including matching-covered graphs and ear decompositions, have been highlighted and used to obtain asymptotically tight approximation results for RAP. The approximation algorithm presented for the general variant of RAP combines classical results for matching-covered graphs with LP randomized rounding techniques.

Some ongoing and future work includes the following lines of research. Study a version of RAP with node failures, or with a combination of node and edge failures. This problem has many potential applications beyond the ones listed here. Study the variant of RAP where each scenario consists of at most k edges, for some input parameter $k > 1$. This paper treats the case $k = 1$. Besides, it is interesting to study the complexity of RAP in general graphs.

Acknowledgement. The work of the second author and the third author is part of the Research Training Group *"Discrete Optimization of Technical Systems under Uncertainty"* (RTG 1855). Financial support through RTG 1855 of the second author by the German Research Foundation (DFG) is gratefully appreciated.

7 References

[1] David Adjiashvili. *Structural Robustness in Combinatorial Optimization.* PhD thesis, ETH Zürich, Zürich, Switzerland, 2012. URL `http://dx.doi.org/10.3929/ethz-a-007579109`.

[2] David Adjiashvili. Non-Uniform Robust Network Design in Planar Graphs. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, volume 40 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 61–77, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL `http://drops.dagstuhl.de/opus/volltexte/2015/5294`.

[3] David Adjiashvili, Sebastian Stiller, and Rico Zenklusen. Bulk-robust combinatorial optimization. *Mathematical Programming*, 149(1):361–390, 2015. URL `http://dx.doi.org/10.1007/s10107-014-0760-6`.

[4] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Complexity of the min–max and min–max regret assignment problems. *Operations Research Letters*, 33(6):634 – 640, 2005. URL `http://www.sciencedirect.com/science/article/pii/S0167637704001658`.

[5] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Min–max and min–max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research*, 197(2):427 – 438, 2009. URL `http://www.sciencedirect.com/science/article/pii/S0377221708007625`.

[6] Robert E. L. Aldred and Michael D. Plummer. On matching extensions with prescribed and proscribed edge sets II. *Discrete Mathematics*, 197–198:29–40, 1999. URL `http://www.sciencedirect.com/science/article/pii/S0012365X99900353`.

[7] Robert E. L. Aldred, Derek A. Holton, M. I. Porteous, and Michael D. Plummer. Two results on matching extensions with prescribed and proscribed edge sets. *Discrete Mathematics*, 206(1–3):35–43, 1999. URL `http://www.sciencedirect.com/science/article/pii/S0012365X98004051`.

[8] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In Giancarlo Bongiovanni, Daniel Pierre Bovet, and Giuseppe Di Battista, editors, *Algorithms and Complexity: Third Italian Conference, CIAC '97 Rome, Italy, March 12–14, 1997 Proceedings*, pages 288–298. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997. URL `http://dx.doi.org/10.1007/3-540-62592-5_80`.

[9] Dimitris Bertsimas, David B. Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53:464–501, 2011. URL `http://dx.doi.org/10.1137/080734510`.

[10] Robert C. Brigham, Frank Harary, Elizabeth C. Violin, and Jay Yellen. Perfect-matching preclusion. *Congressus Numerantium*, 174:185–192, 2005.

[11] Rainer E. Burkard, Mauro Dell'Amico, and Silvano Martello. *Assignment Problems, Revised Reprint:*. Society for Industrial and Applied Mathematics, 2012.

[12] Giuliana Carello and Ettore Lanzarone. A cardinality-constrained robust model for the assignment problem in Home Care services. *European Journal of Operational Research*, 236(2):748–762, 2014. URL `http://www.sciencedirect.com/science/article/pii/S0377221714000113`.

[13] Shiri Chechik and David Peleg. Robust Fault Tolerant Uncapacitated Facility Location. In Jean-Yves Marion and Thomas Schwentick, editors, *27th International Symposium on Theoretical Aspects of Computer Science*, volume 5 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 191–202, Dagstuhl, Germany, 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL `http://drops.dagstuhl.de/opus/volltexte/2010/2454`.

[14] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault-tolerant spanners for general graphs. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 435–444, 2009. URL `http://doi.acm.org/10.1145/1536414.1536475`.

[15] Chandra R. Chegireddy and Horst W. Hamacher. Algorithms for finding k-best perfect matchings. *Discrete Applied Mathematics*, 18(2):155 – 165, 1987. URL http://www.sciencedirect.com/science/article/pii/0166218X87900175.

[16] Chandra Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–129, 2007. URL http://doi.acm.org/10.1145/1324215.1324236.

[17] Eddie Cheng, Linda Lesniak, Marc J. Lipman, and Laszlo Liptak. Conditional matching preclusion sets. *Information Sciences*, 179(8):1092–1101, 2009. URL http://www.sciencedirect.com/science/article/pii/S0020025508004374.

[18] Joseph Cheriyan, András Sebő, and Zoltán Szigeti. Improving on the 1.5-Approximation of a Smallest 2-Edge Connected Spanning Subgraph. *SIAM Journal on Discrete Mathematics*, 14(2):170–180, 2001. URL http://epubs.siam.org/doi/abs/10.1137/S0895480199362071.

[19] Mabel C. Chou, Geoffrey A. Chua, Chung-Piaw Teo, and Huan Zheng. Design for process flexibility: Efficiency of the long chain and sparse structure. *Operations Research*, 58(1):43–58, 2010. URL http://dx.doi.org/10.1287/opre.1080.0664.

[20] Marie-Christine Costa. Persistency in maximum cardinality bipartite matchings. *Operations Research Letters*, 15(3):143 – 149, 1994. URL http://www.sciencedirect.com/science/article/pii/0167637794900493.

[21] Andreas Darmann, Ulrich Pferschy, Joachim Schauer, and Gerhard J. Woeginger. Paths, trees and matchings under disjunctive constraints. *Discrete Applied Mathematics*, 159(16):1726–1735, 2011. URL http://www.sciencedirect.com/science/article/pii/S0166218X1000435X.

[22] Marcelo H. de Carvalho and Joseph Cheriyan. An O(VE) algorithm for ear decompositions of matching-covered graphs. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 415–423, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. URL http://dl.acm.org/citation.cfm?id=1070432.1070490.

[23] Vladimir G. Deineko and Gerhard J. Woeginger. On the robust assignment problem under a fixed number of cost scenarios. *Operations Research Letters*, 34(2):175 – 179, 2006. URL http://www.sciencedirect.com/science/article/pii/S0167637705000465.

[24] Antoine Désir, Vineet Goyal, Yehua Wei, and Jiawei Zhang. Sparse process flexibility designs: Is the long chain really optimal? *Operations Research*, 64(2):416–431, 2016. URL http://dx.doi.org/10.1287/opre.2016.1482.

[25] Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: Better and simpler. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '11, pages 169–178, New York, NY, USA, 2011. ACM. URL http://doi.acm.org/10.1145/1993806.1993830.

[26] Mitre Costa Dourado, Dirk Meierling, Lucia D. Penso, Dieter Rautenbach, Fabio Protti, and Aline Ribeiro de Almeida. Robust recoverable perfect matchings. *Networks*, 66(3):210–213, 2015. URL http://dx.doi.org/10.1002/net.21624.

[27] Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998. URL http://doi.acm.org/10.1145/285055.285059.

[28] Ryo Fujita, Yusuke Kobayashi, and Kazuhisa Makino. Robust matchings and matroid intersections. In Mark de Berg and Ulrich Meyer, editors, *Algorithms – ESA 2010: 18th Annual European Symposium, Liverpool, UK, September 6-8, 2010. Proceedings, Part II*, pages 123–134, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. URL http://dx.doi.org/10.1007/978-3-642-15781-3_11.

[29] Harold N. Gabow, Shachindra N. Maheshwari, and Leon J. Osterweil. On Two Problems in the Generation of Program Test Paths. *IEEE Transactions on Software Engineering*, SE-2(3):227–231, 1976. URL http://dx.doi.org/10.1109/TSE.1976.233819.

[30] Harold N. Gabow, Michel X. Goemans, Éva Tardos, and David P. Williamson. Approximating the smallest K-edge connected spanning subgraph by LP-rounding. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 562–571, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. URL http://dl.acm.org/citation.cfm?id=1070432.1070511.

[31] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. A Series of books in the mathematical sciences. W. H. Freeman, 1979.

[32] Fabrizio Grandoni, R. Ravi, Mohit Singh, and Rico Zenklusen. New approaches to multi-objective optimization. *Mathematical Programming*, 146(1):525–554, 2014. URL http://dx.doi.org/10.1007/s10107-013-0703-7.

[33] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Algorithms and Combinatorics. Springer Berlin Heidelberg, 1993.

[34] MohammadTaghi Hajiaghayi, Nicole Immorlica, and Vahab S. Mirrokni. Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks. In *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking*, MobiCom '03, pages 300–312, New York, NY, USA, 2003. ACM. URL http://doi.acm.org/10.1145/938985.939016.

[35] Refael Hassin and Shlomi Rubinstein. Robust matchings. *SIAM Journal on Discrete Mathematics*, 15(4):530–537, 2002. URL http://dx.doi.org/10.1137/S0895480198332156.

[36] Willy Herroelen and Leus Roel. Project scheduling under uncertainty: Survey and research potentials. *European Journal of Operational Research*, 165(2):289 – 306, 2005. URL http://www.sciencedirect.com/science/article/pii/S0377221704002401.

[37] Chun-Nan Hung, Lih-Hsing Hsu, and Ting-Yi Sung. The most vital edges of matching in a bipartite graph. *Networks*, 23(4):309–313, 1993. URL http://dx.doi.org/10.1002/net.3230230413.

[38] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001. URL http://dx.doi.org/10.1007/s004930170004.

[39] Kamal Jain and Vijay V. Vazirani. An approximation algorithm for the fault tolerant metric facility location problem. In Klaus Jansen and Samir Khuller, editors, *Approximation Algorithms for Combinatorial Optimization: Third International Workshop, APPROX 2000 Saarbrücken, Germany, September 5–8, 2000 Proceedings*, pages 177–182, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg. URL http://dx.doi.org/10.1007/3-540-44436-X_18.

[40] Panos Kouvelis and Gang Yu. *Robust discrete optimization and its applications*, volume 14 of *Nonconvex Optimization and Its Application*. Kluwer Academic Publishers, 1997.

[41] Pierre Laroche, Franc Marchetti, Sébastien Martin, and Zsuzsanna Róka. Bipartite complete matching vertex interdiction problem: Application to robust nurse assignment. In *Control, Decision and Information Technologies (CoDIT), 2014 International Conference on*, pages 182–187, 2014. URL http://dx.doi.org/10.1109/CoDIT.2014.6996890.

[42] László Lovász and Michael D Plummer. *Matching theory*. Volume 121 of North-Holland Mathematics Studies. North-Holland Publishing Co., Amsterdam, 1986.

[43] Neil K. Olver. *Robust network design*. PhD thesis, Department of Mathematics and Statistics, McGill University, Montréal Québec, Canada, 2010.

[44] Temel Öncan, Ruonan Zhang, and Abraham P. Punnen. The minimum cost perfect matching problem with conflict pair constraints. *Computers & Operations Research*, 40(4):920–930, 2013. URL http://www.sciencedirect.com/science/article/pii/S0305054812002481.

[45] Ján Plesník. Connectivity of regular graphs and the existence of 1-factors. *Matematický časopis*, 22(4):310–318, 1972.

[46] M.I. Porteous and Robert E.L. Aldred. Matching extensions with prescribed and forbidden edges. *Australasian Journal Of Combinatorics*, 13: 163–174, 1996.

[47] Jean-Charles Régin. A filtering algorithm for constraints of difference in CSPs. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (Vol. 1)*, AAAI '94, pages 362–367, Menlo Park, CA, USA, 1994. American Association for Artificial Intelligence. URL http://dl.acm.org/citation.cfm?id=199288.178024.

[48] András Sebő and Jens Vygen. Shorter tours by nicer ears: 7/5-approximation for the graph-tsp, 3/2 for the path version, and 4/3 for two-edge-connected subgraphs. *Combinatorica*, 34(5):597–629, 2014. URL http://dx.doi.org/10.1007/s00493-014-2960-3.

[49] David Simchi-Levi and Yehua Wei. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations Research*, 60(5):1125–1141, 2012. URL http://dx.doi.org/10.1287/opre.1120.1081.

[50] David Simchi-Levi and Yehua Wei. Worst-case analysis of process flexibility designs. *Operations Research*, 63(1):166–185, 2015. URL http://dx.doi.org/10.1287/opre.2014.1334.

[51] Chaitanya Swamy and David B. Shmoys. Fault-tolerant facility location. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, pages 735–736, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics. URL http://dl.acm.org/citation.cfm?id=644108.644228.

[52] Christopher S. Tang. Robust strategies for mitigating supply chain disruptions. *International Journal of Logistics Research and Applications*, 9 (1):33–45, 2006. URL http://www.tandfonline.com/doi/abs/10.1080/13675560500405584.

[53] Tamir Tassa. Finding all maximally-matchable edges in a bipartite graph. *Theoretical Computer Science*, 423:50–58, 2012. URL http://www.sciencedirect.com/science/article/pii/S0304397511010474.

[54] Carlos E. Valencia and Marcos C. Vargas. Optimum matchings in weighted bipartite graphs. *Boletín de la Sociedad Matemática Mexicana*, 22(1):1–12, 2016. URL http://link.springer.com/10.1007/s40590-015-0065-7.

[55] Xiumei Wang, Jinjiang Yuan, and Sujing Zhou. Edge-deletable IM-extendable graphs with minimum number of edges. *Discrete Mathematics*, 309(16):5242 – 5247, 2009. URL http://www.sciencedirect.com/science/article/pii/S0012365X09001794.

[56] Gang Yu and Jian Yang. On the robust shortest path problem. *Computers & Operations Research*, 25(6):457 – 468, 1998. URL http://www.sciencedirect.com/science/article/pii/S0305054897000853.

[57] Rico Zenklusen. Matching interdiction. *Discrete Applied Mathematics*, 158 (15):1676 – 1690, 2010. URL http://www.sciencedirect.com/science/article/pii/S0166218X10002143.